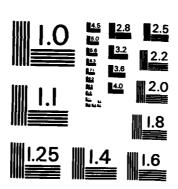
DESIGN PRINCIPLES FOR PERSONALIZED DECISION AIDING: AN APPLICATION TO TAC. (U) DECISION SCIENCE CONSORTIUM INC FALLS CHURCH YA M S COHEN ET AL. JUL 85 DSCI-85-3 RADC-TR-85-140 F36602-84-C-0176 F76 5/8 AD-A159 508 1/2 UNCLASSIFIED NL



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS -1963 - A

RADC-TR-85-140 Final Technical Report July 1985



DESIGN PRINCIPLES FOR PERSONALIZED DECISION AIDING: AN APPLICATION TO TACTICAL AIR FORCE ROUTE PLANNING

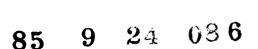
Decision Science Consortium, Inc.

Marvin S. Cohen; Bryan B. Thompson and James O. Chinnis, Jr.

APPROVED FOR PUBLIC RELEASE: DISTRIBUTION UNLIMITED

THE FILE COP

ROME AIR DEVELOPMENT CENTER
Air Force Systems Command
Griffiss Air Force Base, NY 13441-5700



This report has been reviewed by the RADC Public Affairs Office (PA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

RADC-TR-85-140 has been reviewed and is approved for publication.

APPROVED:

Delove P.M. Clark

DELORES P. M. CLARK, 1LT, USAF

Project Engineer

APPROVED:

RAYMOND P. URTZ, JR. Technical Director

Command & Control Division

FOR THE COMMANDER:

RICHARD W. POULIOT

Plans Office

If your address has changed or if you wish to be removed from the RADC mailing list, or if the addressee is no longer employed by your organization, please notify RADC (COAD) Griffiss AFB NY 13441-5700. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document requires that it be returned.

1a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED 2a. SECURITY CLASSIFICATION AUTHORITY N/A 2b. DECLASSIFICATION / DOWNGRADING SCHEDULE N/A Approved for public release; distribution unlimited.					
2a. SECURITY CLASSIFICATION AUTHORITY N/A 2b. DECLASSIFICATION / DOWNGRADING SCHEDULE 3. DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release; distribution	Į.				
N/A 2b. DECLASSIFICATION / DOWNGRADING SCHEDULE Approved for public release; distribution					
N/A Unlimited.	i				
4. PERFORMING ORGANIZATION REPORT NUMBER(S) 5. MONITORING ORGANIZATION REPORT NUMBER(S)	ſ				
85-3 RADC-TR-85-140	RADC-TR-85-140				
63. NAME OF PERFORMING ORGANIZATION 65. OFFICE SYMBOL 73. NAME OF MONITORING ORGANIZATION					
Decision Science Consortium, Inc. (If applicable) Rome Air Development Center (COAD)	Rome Air Development Center (COAD)				
6c. ADDRESS (City, State, and ZIP Code) 7b. ADDRESS (City, State, and ZIP Code)	7b. ADDRESS (City, State, and ZIP Code)				
7700 Leesburg Pike					
Suite 421 Griffiss AFB NY 13441-5700	Griffiss AFB NY 13441-5700				
Falls Church VA 22043					
8a. NAME OF FUNDING/SPONSORING 8b. OFFICE SYMBOL ORGANIZATION (If applicable) 9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER					
Rome Air Development Center COAD F30602-84-C-0176	ł				
8c. ADDRESS (City, State, and ZIP Code) 10. SOURCE OF FUNDING NUMBERS					
PROGRAM PROJECT TASK WORK UN	ıΤ				
Griffiss AFB NY 13441-5700 ELEMENT NO. NO. NO. ACCESSION 65502F 3005 RA 01	NO.				
11 TITLE (Include Security Classification) DESIGN PRINCIPLES FOR PERSONALIZED DECISION AIDING: An Application to Tactical Air Force Route Planning					
12. PERSONAL AUTHOR(S)-					
Marvin S. Cohen, Bryan B. Thompson, James O. Chinnis, Jr.	- 1				
13a TYPE OF REPORT 13b. TIME COVERED 14. DATE OF REPORT (Year, Month, Day) 15. PAGE COUNT Final FROM Sep 84 to Mar 85 July 1985					
16. SUPPLEMENTARY NOTATION					
N/A	Ì				
17 COSATI CODES 18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)					
FIELD GROUP SUB-GROUP Decision Aids Displays Route Planning	1				
05 08 Tactical Air Force Human Engineering					
09 04 Expert Systems Cognitive Processes					
19. ABSTRACT (Continue on reverse if necessary and identify by block number)					
A major problem in the design of decision aids is the combination of flexibility and	1				
effective normative guidance in procedures for representing, manipulating, and displaying	;				
uncertainty. The objectives of the present research are: (a) to develop a consistent	ı				
framework for personalized decision aiding; and (b) to apply that framework to a specific					
RADC decision aid. A framework has been developed, based on an analysis of decision mak	.ng				
into cognitive subtasks, and on the specification of generic interface functions for	ļ				
personalizing, channeling, and advisory prompting. This framework was applied in an					
analysis of Tactical Air Force route planning and in a critique of a current decision					
aid in that context (RPA). A Personalized Route Planner was designed and a demonstration					
system has been implemented to illustrate and test the feasibility of principles of					
personalized decision aiding. Alived Courses to the transfer of the second to the seco					
from fortage inquiries					
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT 21. ABSTRACT SECURITY CLASSIFICATION					
UNCLASSIFIED UNCLASSIFIED					
22a. NAME OF RESPONSIBLE INDIVIDUAL 22b. TELEPHONE (Include Area Code) 22c. OFFICE SYMBOL					

ACKNOWLEDGMENTS

The work reported here was sponsored by the Rome Air Development Center under a Small Business Innovative Research (Phase I) contract. We are grateful to our Technical Monitor, LT Delores P.M. Clark for her help and guidance. We are also grateful to Professor John Payne of Duke University for significant contributions and clarifications of ideas; and to Edger Stephenson (Colonel, USAF ret.) and Stephen Matthews (Captain, USAF ret.) for sharing their knowledge of the route planning environment and reviewing our ideas for the Personalized Route Planner.

::C:	on Fur	
DTIC	ounced	
By Dict ib:	ution/	
A	vailability C	odes -
Oist A.(Avail and Special	



TABLE OF CONTENTS

			Page
1.0	INTROI	DUCTION	1
	1.1 1.2 1.3	The Problem	1 1 3
2.0	A FRAN	MEWORK FOR PERSONALIZED DECISION AIDING	5
	2.1 2.2 2.3 2.4 2.5	Importance of Pre-Existing Decision Processes Current Technology	6 8 10 12 12 14 18 19 19 20 21 21 22 25 27
		2.6.4.3 Imprecision	
3.0	ROUTE	PLANNING	33
	3.1 3.2 3.3	The Decision Task	35

TABLE OF CONTENTS (Continued)

		I	?age
4.0	A PEI	RSONALIZED ROUTE PLANNING AID	39
	4.1 4.2 4.3 4.4 4.5 4.6 4.7 4.8 4.9	Implementation Overview of the User Interface Overview of Menu Organization CONSTRAINTS FOCUS TOPIC MODIFY GENERATE ANALYZE	39 40 42 43 45 46 50 52 55
5.0	CONC	LUSIONS	63
	5.1 5.2 5.3	Feasibility of Personalized Decision Aiding Prescriptive Mechanisms for PRP	63 63 68
REFER	ENCES		73
APPEN	DIX:	DISPLAY SCREENS	81
PROGRA	AM LIS	STING	VOL II

1.0 INTRODUCTION

1.1 The Problem

In Air Force combat environments, achievement of command and control objectives depends not only on effective data collection and communication, but also on processes of problem-structuring, inference, and decision after the data have been collected and transmitted. A major goal of the Rome Air Development Center, therefore, is to explore ways in which computerized support can be provided for the rapid integration of battlefield data into higher-level processes of reasoning and choice.

A pervasive aspect of those processes is the handling of uncertainty. For example, in tactical Air Force intelligence analysis, the identities, locations, capabilities, and intentions of threats must be inferred from data that are typically incomplete, unreliable, inconsistent, and constantly changing. In operational analysis, activities such as target selection, route planning, and allocation of resources to missions involve a crucial and often subtle balancing of tradeoffs among risks to friendly assets and uncertain opportunities against targets. The representation, manipulation, and display of uncertainty is central both for the decision maker's understanding of his own problem and for the consistent exchange of information among diverse decision makers and among different automated systems.

1.2 Technical Challenges

Advances in decision theory and in artificial intelligence have led to a heightened interest in aids to support inference (e.g., Edwards, 1966; Brown et al., 1974; Cohen and Brown, 1980). Such a development is not surprising. Although uncertainty and the need to act in the face of uncertainty are pervasive, especially in warfare, there is growing evidence in the field of cognitive psychology that humans are less than optimal in

probabilistic reasoning (e.g., Kahneman, Slovic, and Tversky, 1982). Yet neither the identification of a need nor the provision of a solution has proven as straightforward as originally anticipated.

Tactical Air Force analytical problems—as in other cases of reasoning with uncertainty—tend to mix intuition and analysis in a multiplicity of ways. At times, an appropriate conclusion may be immediately apparent; at other times, raw data must be carefully interpreted and integrated, intermediate conclusions or hypotheses may be derived, alternative accounts or options may need to be generated, and competing lines of argument weighed. There is potential variability among decision makers and decision problems at virtually every step in this process, from initial identification of the problem to reporting a solution. Cognitive psychologists are only now beginning to explore the dimensions of this variability, but there is evidence that failure to accommodate it in a computer-based aid may lead to rejection and disuse (e.g., Miller, 1980; Beard, 1977; Cohen et al., 1982).

There is, in short, a conflict in decision aid design between prescriptive and descriptive concerns: on the one hand, catering to the preferences of the user, whatever they may be; on the other hand, steering users away from pitfalls and fallacies to which a preferred decision strategy might give rise. Because of this conflict, we suspect, high-level users of computer-based information systems typically find that either too little or too much help is offered. Users are caught between systems that serve as passive tools (i.e., which automate routine functions like data storage, sorting, and retrieval) and systems which tend to dominate any dialogue with the decision maker (e.g., decision aids and expert systems).

The problem is exacerbated, ironically, by ambiguity concerning the "correct" prescriptive approach. There is increasing debate among statisticians and logicians regarding the relevance of diverse concepts of uncertainty (e.g., chance, incompleteness of evidence, and imprecision) and the appropriate formal calculi for representing them (cf., Shafer, 1976; Zadeh,

1965; Cohen, Watson, and Barrett, 1985; Cohen, Schum, Freeling, and Chinnis, 1984). Some theorists (e.g., L.J. Cohen, 1981) have even argued that where prescriptive standards and human practice diverge, it is often the prescriptive theories that are at fault rather than practice.

What, then, is the role--if any--for effective and normatively defensible decision guidance? Can it coexist, within a decision aid, with knowledge structures and information-processing strategies that are compatible with those of users?

1.3 Outline of the Research

In the present work, we have addressed these questions in a specific tactical Air Force operational environment. The objectives have been: (1) to develop and demonstrate principles of information display and decision aid design that are both personalized and prescriptively valid; and (2) to test the feasibility of those principles in a specific RADC decision aid.

A variety of aids developed at RADC were examined: in particular, the Target Prioritization Aid (TPA), Duplex Army Radio/Radar Targeting Aid (DART), Route Planning Aid (RPA), and Dynamic Air Order of Battle Aid (DAGR). All of these aids address, in different ways, the problem of weighing uncertainty. Moreover, they apply to a set of closely interrelated decision-making tasks in the Tactical Air Control Center (TACC) and elsewhere. Outputs of one aid would, ideally, serve as inputs to another. Unfortunately, in their present forms, the aids differ in their conceptualizations of uncertainty, in their ways of computing with it, and in their manner of communicating it to users. An additional objective of the present work, therefore, was to address the need for an integrated cognitive interface for this set of decision aids.

Design of a personalized and prescriptive aid requires a close look at the decision-making environment, on the one hand, and the cognitive processes

of potential aid users in that environment, on the other. For this reason, the most productive initial line of attack was an in-depth look at a single aid, keeping in mind, however, the eventual requirement for a unified framework. The Route Planning Aid was finally chosen as a testbed for the initial phase of this research, although the concepts and methods that have been developed promise to have a far wider applicability.

The results of the present work fall into two categories: initial steps in the development of a systematic framework for personalized decision aid design, and the application of that framework to route planning. This framework, which is described in Section 2.0, is based on findings and theories in cognitive psychology, and on applied experience in the design of decision aids. Section 3.0 uses the framework to analyze the route planning task and the current Route Planning Aid. In Section 4.0 we present the design of a Personalized Route Planner which incorporates the basic principles of personalized decision aiding. The implementation and demonstration of this system has proven, we believe, the feasibility and promise of the basic concepts involved. In Section 5.0 we turn to conclusions and future directions.

2.0 A FRAMEWORK FOR PERSONALIZED DECISION AIDING

2.1 Importance of Pre-Existing Decision Processes

To what extent does it matter, in designing a decision aid, that we understand the cognitive mechanisms involved in unaided decision making? The introduction of decision aids is usually justified on the basis of time (speeding up the processing of large quantities of data) and quality (e.g., avoiding biases and fallacies in unaided reasoning). In neither case has there been much impetus to utilize a detailed understanding of pre-aided knowledge representations and processes of reasoning. The logic of aid design often appears to be: (1) that the pre-aided approach is too slow and/or biased (i.e., different from some favored prescriptive theory) -hence, (2) it needs to be replaced in toto. Decision aid concepts and inferential processes, therefore, need bear no resemblance to the concepts and processes existing before the advent of the aid. A somewhat more sophisticated rationale for ignoring pre-existing decision-making practice is that human reasoning is itself highly adaptive; there is no "natural" approach. New methods of thinking will simply come into being after users learn to use faster and/or prescriptively correct decision-aiding tools.

In our view, the balancing of descriptive and prescriptive concerns is not so straightforward. Unaided decision-making processes may sometimes be biased, but they may also reflect expertise accumulated over years of experience in a domain. Current decision-aiding technology is often incapable of replicating the full extent of such expertise. Buchanan (1981) and McCarthy (1977) list a variety of concepts for which current artificial intelligence methods are, at least in part, inadequate. These include causal reasoning, propositional attitudes like intention or belief, conflicting plans, analogical reasoning, and reasoning about dynamic three-dimensional relationships. Computers are not (yet) as good as humans at reasoning on multiple levels, generating novel solution methods when current approaches fail (cf., Newell, 1981), and handling unanticipated types

of data (cf., Johnson, 1985). Paradoxically, then, decision aids may diminish the overall quality of performance if they prevent users from thinking in their usual ways, and if they fail to offer methods for easily incorporating user insights into automated processes of reasoning.

2.2 Current Technology

Among the major sources of decision-aiding technology are operations research, decision analysis, and artificial intelligence. These can be arranged along a rough continuum by the degree to which they model actual human decision processes rather than mathematically derived relationships. Nevertheless, none of them has traditionally focused on the potential for collaborative problem-solving between users and computers:

- (1) Operations research: Typically only a subset of the issues in a given problem can be modeled by operations research techniques: those amenable to objective measurement and rigorously defined mathematical optimization (cf., Watson, 1981). Since actual decisions almost always involve "soft" factors, users are left with the choice of disregarding the output of the model, ignoring their own intuitions, or--somehow--integrating the outputs of such models with their personal methods of conceptualizing and reasoning.
- (2) Decision analysis: In decision analysis, the focus shifts from the problem to the decision maker. Subjective judgments regarding values and neertainties, as well as objective data, are incorporated within an axiomatically justified calculus. The result, however, is an "idealized" version of the decision maker that may not effectively tap what the real decision maker knows. In computer-based aids, decision analysis often imposes a rather rigid allocation of cognitive tasks: users are assumed to have precise and valid representations of components of the problem (e.g., likelihood ratios which quantify the impact of specific items of evidence), but not of the whole problem. The computer has the task of aggregating the various inputs from the user (and from other sources). It is more

Advisory Prompts for Select:

- Research in cognitive psychology suggests that humans tend to seek additional confirming evidence for a favored hypothesis. A computer aid might monitor a user's pattern of information requests, examine its own model of the problem in order to draw inferences about the hypotheses the user has in mind, and prompt the user if evidence or hypotheses exist which the user has failed to consider but which may have an impact on his conclusions.
- Humans often find it difficult to assess the overall credibility of a conclusion based on several steps of reasoning; they simplify by ignoring the uncertainty at early stages. Prompts might warn users, when they appear to be acting as if a particular hypothesis were known to be true, that a number of stages of uncertainty must be kept in mind. The same type of caution might be appropriate when a compound, or conjunctive, hypothesis is being considered.
- The user might be notified when two information sources, both of which are regarded as credible, have contradicted one another. He might then choose to readjust one or both credibility assessments downward. An Advisory prompt might notify him on future occasions when either of the (partially) discredited sources is involved in an important conclusion.
- 2.6.1.2 Modify. While the aid should permit user adjustment of any meaningful values employed in the database, it should selectively facilitate adjustment of values about which users are likely to have information not available to the computer. Values to be adjusted should be decomposed into parameters about which users are likely to have reliable intuitions. Automatically computed values should be displayed as a reference, so users can focus on the appropriate direction and magnitude of the adjustment (based on the new evidence) and not have to integrate all the evidence to come up with an absolute value.

Advisory Prompts for Modify:

• Humans tend to combine evidence by a process that is more like averaging than like proper Bayesian inference. When adjustments fit an averaging pattern, prompts might remind subjects to

only sketch out the shape some such guidelines might take. We first identify candidate channeling and prompting functions corresponding to the five interface functions of Section 2.4, and designed to address some of the psychological research findings outlined in Section 2.5. Other potential guidelines are then discussed under several headings: allocation of cognitive tasks, channeling versus prompting, and the role of alternative prescriptive concepts.

- 2.6.1 Generic Interface Functions for Channeling and Prompting. There has as yet been little empirical research demonstrating the efficacy of "debiasing" techniques. Nevertheless, in cases where debiasing has been successful (e.g., Lopes, 1982; Slovic and Fischhoff, 1977; Koriat, Lichtenstein, and Fischhoff, 1980), there appears to be a common element. In these studies improved performance was achieved not by directly instructing subjects in the prescriptively correct approach, but by methods more attuned to natural psychological processes. We now outline some potential procedures of this type, corresponding to the generic interface functions of Section 2.4 and utilizing principles of channeling and advisory prompting. (In each case, channeling functions are discussed first, followed by concepts for advisory prompting.)
- 2.6.1.1 Select: (Focus) While users should be able to organize displays around a variety of meaningful user-designated objects, the aid should facilitate the use of decision-related objects for this purpose. For example, the aid can facilitate clustering of options by their performance on a selected evaluative criterion, or help isolate components of an option (e.g., segments of a route) whose design is affected by a specific factor (e.g., a particular threat). (Topic) When an intermediate result or conclusion is uncertain, the sources of its uncertainty should be explicitly indicated. Evidence for a result should be available for display along with the result. Inferential relationships in the database can be "mapped" by menus, which permit tracing a process of reasoning from its sources of evidence to its final conclusion.

on an attribute are eliminated at each stage, and not considered further. In this strategy, an option might be eliminated for missing a cut-point on one dimension even though it scores very high on other dimensions. Tradeoffs, or compensatory relations among dimensions, are thus not considered. In another heuristic strategy called "satisficing" (Simon, 1957; Svenson, 1979), information search is organized by options. The decision maker considers a sequence of options until he finds one that clears the cut-points he has selected on relevant attributes. Here again compensatory relationships are ignored. Payne (1981) has suggested that these diverse information search strategies may correspond to different ways in which decision makers organize knowledge.

2.6 Principles for Personalized Aid Design

The design of channeling and prompting functions demands a precarious balancing act involving: (a) the desirability of maximizing flexibility and personalization; (b) the apparent likelihood, based on laboratory data, of biases and fallacies in reasoning; and (c) the potential applicability of multiple concepts of uncertainty in the interpretation of natural processes of reasoning. Cognitive psychologists have now begun to suggest corrective procedures for biases and fallacies that take account of preferred methods of decision making (Kahneman and Tversky, 1979; Fischhoff, 1982; Lopes, 1982). But there has been little progress as yet on the human engineering of these "debiasing" procedures in real-world contexts. At the same time, some theorists have argued that suboptimal cognitive performance may, in part at least, be a function of: (a) the artificiality of task demands (Neisser, 1976; Hogarth, 1981); or (b) a misunderstanding of the concepts of uncertainty which humans in fact seek to apply (L.J. Cohen, 1981). Yet little has been done thus far to implement such insights in aids that match natural human concepts.

To what extent can these tradeoffs be reconciled within a coherent set of guidelines for personalized decision aid design? At this stage, we can

fects will depend on the degree to which decision makers lack confidence in the probability estimates. This, in turn, may depend on the degree to which evidence for an estimate matches the type of evidence represented in user knowledge structures. An additional set of biases involves distorted conceptions of randomness in everyday judgment, e.g., the "gambler's fallacy" where a sequence of similar outcomes, which are in fact independent, is thought to increase the likelihood of a different outcome on the next trial. Fallacies of this sort may be inevitable by-products of powerful top-down or expectancy-driven processes of pattern recognition (Lopes, 1982).

Assess Value of Outcomes. Decision makers do not typically consider all the potential outcomes of an action together. Rather, outcomes are grouped into "mental accounts" corresponding to natural objects or causal relations, and choices may depend critically on the particular grouping that is adopted (Kahneman and Tversky, 1982). An additional cognitive simplification is achieved by representing an outcome in causally relevant terms, by the difference it would make relative to some reference point. Decisions may be significantly affected by the choice of reference levels, since the same outcome may be regarded as a gain or as a loss. For example, the outcome of a defensive tactic may be encoded as 400 men saved (relative to the number who would have died had nothing been done) or as 200 men lost (relative to the status quo). An important finding by Kahneman and Tversky (1979) is that decision makers are more likely to take risks when outcomes are represented as losses than when they are represented as gains.

Select an Option. Heuristic procedures may be adopted which reduce the cognitive effort that would be required in a thorough consideration of every option. Such heuristics have implications for the way decision makers search information. In Elimination by Aspects (Tversky, 1972), for example, search is organized by evaluative attributes. Attributes are considered serially in order of importance; options falling below a cut-point

process by acting as if conclusions at earlier stages (e.g., range) were known to be true, rather than merely inferred (Schum, Du Charme, and DePitts, 1973). Similarly, the probability of a detailed hypothesis or scenario is likely to be judged higher than the probabilities for its components (Tversky and Kahneman, 1983). The latter effect may arise because additional details increase the match between the hypothesis and user knowledge structures (Leddo, Abelson, and Gross, 1984).

Option Generation. People segment complex options into "natural" components, and treat the elements as if they were independent choices, leading to suboptimal portfolios (Tversky and Kahneman, 1981). There is a tendency to formulate options in terms of immediate actions that span only a short timeframe rather than as long-term policies, and to overlook, as a result, the cumulative risk of pursuing a given course of action over a long period of time (Slovic, Fischhoff, and Lichtenstein, 1978). Individuals differ in the degree to which they consider future choices in current planning (Streufert and Streufert, 1981) and in the number of options they generate (Driver and Mock, 1976). Ingrained ways of viewing a problem tend to hinder the generation of novel and creative solutions (Pitz, Sachs, and Heerboth, 1980).

Generate Possible Outcomes of Options. In considering what might happen if a particular option is adopted, people are subject to biases based on their internal causal models, as well as biases in the use of past experience, such as a heightened tendency to remember salient events or events that occurred very late or very early in a sequence.

Assess Uncertainty of Outcomes. Some of the biases which affect situation assessment may also occur when predictions are made contingent on a particular option. Additional pitfalls, however, include the effects of "wishful thinking" (e.g., higher probability assessments for high utility outcomes) or overcautiousness (e.g., lower assessments for high utility outcomes). According to Einhorn and Hogarth (1984), the size of these ef-

or consider evidence that bears on other issues (Shaklee and Fischhoff, 1982).

Infer Conclusions. A number of studies, which show that a statistical model of a person's judgment process can outperform (in accuracy) that person's own judgments, suggest that people do not effectively utilize the information available to them in inference tasks (Dawes, 1975; Cohen, 1982). Other laboratory results suggest possible causes. For example, people tend to ignore later evidence that contradicts a favored, or earlier, datum and to double count redundant evidence (Schum and Martin, 1981). Also, people commonly ignore statistical, or "base rate," data and overweight unique or problem-specific factors (Kahneman and Tversky, 1972). Both of these observations suggest the predominance in natural reasoning of non-statistical, causal models (Johnson, 1985). A related line of research suggests that people make predictions about a system by "running" a mental model of that system in their heads, rather than by applying statistical or logical rules (Gentner and Stevens, 1983). Results can be distorted, and overconfidence can occur, when false analogies between the system and the model influence conclusions. When people do attempt to make statistical judgments, moreover, estimates may be biased by the ease of recall (or "availability") of a particular class of events in a mental sampling (Tversky and Kahneman, 1972).

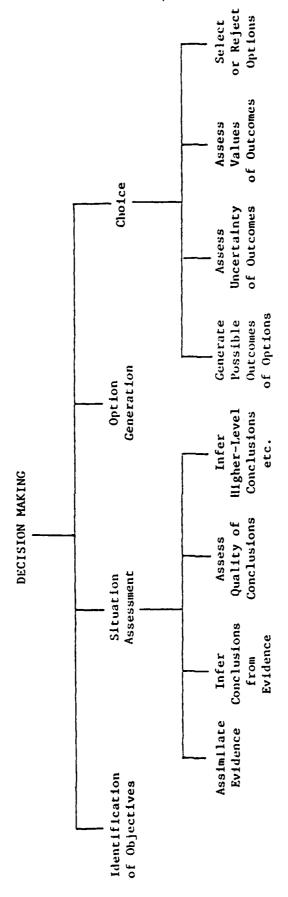
Assess Quality of Conclusions. A number of studies show that people consistently overestimate their degree of certainty regarding predicted events and estimated quantities, even in areas where they are (rightfully) regarded as experts. While there is some evidence that experts (as opposed to college sophomores) are less susceptible to overconfidence (Lichtenstein, Fischhoff, and Phillips, 1982), other research indicates that the difference between expert and novice is slight (Kadane and Lichtenstein, 1982). When inference proceeds in stages (e.g., deriving the probability of being hit by enemy fire from information about the range of a threat, which is derived from bearings data), people often simplify the

collecting and viewing data or evidence, deriving inferences, developing some sense of confidence in the conclusions, and continuing, perhaps, to draw further higher-level inferences. Again, the steps may be iterative, may be combined, or may be skipped altogether by some decision makers in some situations. (Note that the term "evidence" is quite relative; evidence in one process may be the highly uncertain conclusion of a prior analysis.) This decomposition of cognitive subtasks could, of course, be continued. It has been postulated that all cognitive functioning can ultimately be analyzed into a set of simple "elementary information processes" (Newell and Simon, 1972; Chase, 1978) such as selecting an input, reading the value of a variable, comparing two values, and eliminating an alternative.

2.5.2 Cognitive Shortcomings. Each of the cognitive subtasks identified in Figure 2 has been associated, at least in laboratory research, with characteristic shortcomings in reasoning. Thus, by placing recent findings in cognitive psychology within this framework, we may derive a tentative specification of cognitive operations that might benefit from channeling or advisory prompting.

The following outline is highly incomplete and is only meant to touch on some of the issues that bear directly on the present work. Three important themes, however, should emerge: (1) Unaided decision processes employ simplifying heuristics that at best only approximate prescriptively accepted rules (e.g., Bayesian probability theory); (2) a typical effect of such heuristics is that awareness of uncertainty is suppressed; and (3) in many instances, biases are a result of (otherwise successful) efforts to utilize natural knowledge structures and processes of reasoning.

Assimilate Evidence. Patterns of information search in laboratory tasks tend to avoid stringent tests of favored hypotheses (Wason, 1960, 1981; Einhorn, 1980). At the same time, there is a tendency to seek confirming evidence of an already well-supported hypothesis, rather than take action



Potential cognitive subtasks in the decision-making process. Figure 2.

2.5 The Need for a Prescriptive Counterbalance

These five functions, as they stand, provide a sort of informal checklist for decision aid design, as well as a partial framework for the evaluation of a decision aid or decision aid design concept. Unfortunately, however, personalization to this degree may lead to a highly "user-friendly" system with seriously suboptimal outputs. Flexibility must be tempered by the impact of a prescriptive model, through channeling and advisory prompting. The basis for channeling and prompting functions is an understanding of how and when human decision strategies diverge from prescriptively justified approaches.

2.5.1 Cognitive Task Analysis. The decision-making process can be conceptualized quite generally as consisting of a specific set of cognitive tasks (Figure 2). First, goals or objectives must be known or identified (if these are not present, there is no motivation to decide or act). Secondly, current circumstances, insofar as they are relevant to the achievement of a goal, are assessed. If a discrepancy is perceived between goals and reality, options for action may be generated. If more than one option is available, a choice will be made.

This is by no means a rigid sequence: the process can be iterative (for example, revising goals, reassessing the situation, or generating new options when the choice process fails to turn up an acceptable alternative); and steps may be skipped (when, for example, the appropriate action is known based on past experience with very similar situations). But the basic set of possibilities is as shown, at least in most decision contexts, and some such framework is critical, we believe, for identifying the specific aspects of human performance where prescriptive aiding may be of use.

It is convenient to break each of these major tasks down into more specialized cognitive subtasks. For example, situation assessment consists of

where descriptive and prescriptive concerns need to be balanced. In particular, these functions specify how personalization, channeling, and prompting might be accomplished in a set of diverse aiding contexts.

Select: (Focus) Users may personalize displays of information by organizing them around alternative meaningful user-designated objects (e.g., regions, threats, routes). By implication, the knowledge base of the aid must match (in critical respects) user cognitive structures that are associated with inferences or decisions of the type in question. (Topic) The user can examine (preferably through interactive graphics) any significant input, inference rule, intermediate conclusion, or final result in the data base.

Modify: The user can alter values of any data base element and immediately observe the impact on results downstream in a chain of reasoning; users may undo their modifications and restore the original values; user inputs may be at any level of fuzziness or precision.

Generate: Users may define options at any level of detail or generality, at any level of fuzziness or precision, and with respect to any time horizon; complex options may be specified all at once or one component at a time.

Analyze: In the evaluation of an option set, users may organize their exploration of the data base according to any preferred search scheme, corresponding to a variety of heuristic choice strategies; users determine the relative importance of different evaluative criteria.

Alert: The system prompts a user when events occur or facts are learned which would play a significant role in user-preferred modes of reasoning and organizing information.

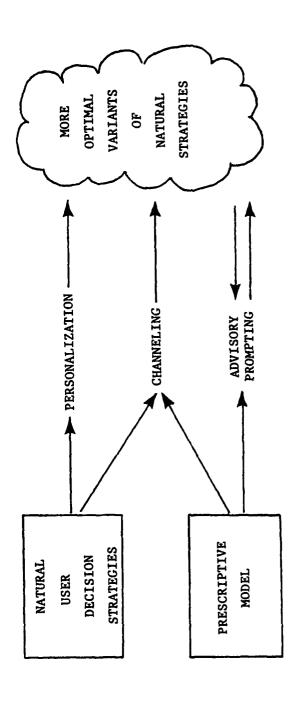
performed tasks and prompts for the addition of procedures that ameliorate potential biases and fallacies; or monitors computer-performed tasks and prompts where human contributions might improve results. Prompting enables the computer to sense weaknesses in a line of reasoning, whether its own or the user's, and to call for (or offer) help.

2.4 Generic Interface Functions for Personalization

Personalization, channeling, and prompting have been utilized in one form or another in a relatively small number of aids. Questions of how much aiding to offer and in what form have been settled on a case-by-case basis, with little "top-down" guidance. To what extent can general principles be developed to guide the detailed application of these three design concepts? Such principles might form the cornerstone of an effective technology for personalized decision aiding. In this and the following sections, we turn to a preliminary examination of that question.

Note that the concept of personalization employed here is one in which the user can actively customize aid functions. The aid does not (in general) diagnose the personality type or "cognitive style" of the user and provide an automatic customization to that type. Experimental data show very little invariance of cognitive style measures across different tasks; i.e., the same user is likely to adopt quite different approaches in different contexts (e.g., Hamm, 1983; Huber, 1982). By contrast, the approach pursued here is one of "efficient flexibility" (Cohen et al., 1982): decision aids are pretuned to facilitate those processing strategies which users are most likely to prefer.

Based on some concepts originally developed for the Navy in a submarine command and control environment (Cohen et al., 1982), we have identified five generic cognitive interface functions. Each enables the user to personalize his interaction with the aid in a different way. They promise, as a set, to be applicable in varying degrees across a wide variety of aids



Three principles for blending descriptive and prescriptive concerns in decision aiding. Figure 1.

2.3 A New Departure

What is lacking is a framework for collaborative problem-solving at the cognitive level. Decision aid performance and usability may both be increased by techniques for effectively and acceptably blending user and computer expertise. How then can we deal with conflicting prescriptive and descriptive objectives? Three design concepts appear relevant (Figure 1):

- Personalizing: Aids facilitate user-preferred representations and information-processing strategies;
- Channeling: Aids provide a context which favors more optimal variants of user-preferred strategies;
- Advisory prompting: Aids prompt for user actions which mesh with, and remedy shortcomings in, user-preferred strategies.

The first principle maximizes the tailoring of person-computer interactions to the particular style of a user. The second and third principles provide a prescriptive counterbalance. Channeling and prompting are designed to prevent or compensate for deviations from optimality that may emerge from personalization, and to do so in the most non-obtrusive way possible. Users, for example, need not grapple directly with the details of a rigorous prescriptive model. Variables and relationships which are important in the prescriptive model, however, are highlighted in displays and explanations to the user. At the same time, the prescriptive model itself functions in the background as the source of occasional "intelligent" advice.

The distinction between channeling and prompting is one of tactics. Channeling encourages the adoption by users of variants of their preferred strategies which are less susceptible to biases and fallacies, by structuring the problem in such a way that those variants become natural and simple to execute. Channeling, therefore, is implicit and proactive. By contrast, prompting is explicit and reactive: the system monitors human-

plausible to suppose, however, that experienced users often have multiple representations at different levels of the same problem (e.g., intuitions about the impact of specific evidence as well as direct intuitions about the overall conclusion). For example, research on expert and novice problem-solving (e.g., Larkin et al., 1980; Chi et al., 1981) suggests that expertise is characterized by an ability to recognize complex patterns and link them directly to appropriate responses, while a sophisticated non-expert may apply a more analytic approach. Perfect confidence or precision may not be warranted at any of these levels, but none should be dismissed. Each may capture some non-overlapping segment of the knowledge which users bring to bear on the problem (cf., Cohen, Schum, Freeling, and Chinnis, 1984).

(3) Knowledge-based expert systems: Expert systems set out to "capture" the actual knowledge of domain experts within a computer. Nevertheless, the power of an expert system model depends, in large part, on choices made by the computer scientist with regard to forms of representation and reasoning (e.g., rule-base, frames, predicate calculus). In many, if not most cases, choices of methods for handling uncertainty have been highly ad hoc. As a result, they lack normative justification, and may in fact lead to highly counterintuitive results in some applications (cf., Buchanan and Shortliffe, 1984). The focus on modeling experts in this tradition has had another consequence: relatively little attention has been given to expert system users. Artificial intelligence contributions to the human-computer interface have, for the most part, focused on input-output issues (e.g., spatial data management, natural language understanding, voice data entry), rather than the design of knowledge representations and inference mechanisms that conform to user requirements. Work in expert systems on explanation and mixed-initiative dialogues has emphasized an essentially passive role for users, as initiators of queries, recipients of answers, and providers of raw, undigested data.

- consider what conclusion a new bit of evidence favors, before performing an adjustment (Lopes, 1982).
- Users should be prompted when information they possess may be of significant value: i.e., when: (1) there is incompleteness of evidence or a conflict among lines of reasoning in the computer model of the problem; (2) the user has potential access to relevant information; and (3) the result is expected to have an impact on choices among actions and ultimate payoffs.
- 2.6.1.3. Generate. Displays should facilitate relatively long time horizons for planning (e.g., by appropriate scaling). Simultaneous specification of all components of a complex option should be supported. Displays should permit generation and simultaneous comparison of multiple options. The aid should facilitate generation of options which include future choices or contingencies.

Advisory Prompts for Generate:

- Short-range planning might be more appropriate in some situations (e.g., where feedback is continuous and mistakes can be easily and quickly corrected), while long-range planning would be more suitable in others (e.g., where a risk appears small unless it is considered cumulatively over the long run). Prompts might recommend that the user consider a shift in the time horizon under appropriate circumstances.
- Users should be prompted if they have generated and evaluated a complex option piece-by-piece and if overall optimality would be significantly improved by considering the option as a whole.
- The user should be prompted if only one option has been generated, but another option exists which is superior on at least one dimension.
- The user should be prompted if contingency plans have not been incorporated in an option, but significant information is likely to become available during its execution.
- 2.6.1.4 Analyze. The user's attention should be drawn to tradeoffs between different evaluative dimensions by displaying scores for an option on

more than one dimension concurrently (e.g., costs and benefits). Action recommendations should be explained by itemizing how options differ on all significant dimensions. The aid should allow the encoding of outcomes in terms of more than one reference point (e.g., assets lost, assets saved).

Advisory prompts for Analyze:

- Humans tend to employ simplified choice schemes that disregard tradeoffs. The user should be notified when he has eliminated an option because it fails to achieve a specified level on a favored evaluative dimension, if that option has significant advantages on other dimensions. The user might be told how much stretching of his specified criterion is required to readmit the rejected option.
- An Advisory prompt should occur when a user entertains an option which is dominated (inferior or tied on all dimensions) by other options.
- 2.6.1.5 Alert. While users are free to designate any item or variable in the database as a criterion for alerting, alerts should also occur on a prescriptive basis. Users should be notified when events occur or facts are learned which have high impact within an appropriate prescriptive model: e.g., which disconfirm previously well-supported inferential hypotheses or which significantly affect choices among actions.
- 2.6.2 Advisory Prompting in Cognitive Task Allocation. Traditionally, task allocation in human-machine systems has been according to the purported strengths of each (e.g., Fitts, 1951). Such methods have for the most part produced a fixed allocation of broadly defined activities, e.g., assigning numerical computation and long-term data storage to the computer and option generation to the human. As noted in Chinnis et al. (1984) and Cohen et al. (1982), this approach fails in application to computer-assisted reasoning. It is not fine-grained enough: variations in task demands and in decision-maker expertise are not captured. It is too machine-oriented: resulting task assignments may not form a meaningful or

organizationally acceptable pattern for human users; users may be unprepared to take over in case of machine dysfunction. Finally, it is too conservative: novel approaches based on human-computer collaboration (e.g., channeling and prompting) may be overlooked.

It is worth noting that the decision-analytic philosophy of "divide and conquer" is in part an application of this rigid task allocation approach: humans assess components of a problem and the machine aggregates them (e.g., Edwards, 1966). The problems are: (1) human assessments of the components may themselves be biased (Section 2.5); (2) humans may possess valid intuitions about the whole problem (indeed, there is an arbitrariness about distinguishing wholes and parts, since the conclusion of one problem can be part of another); and (3) in high-level decisions, there is often organizational resistance against turning responsibility for conclusions over to a computer.

If static generalizations regarding human-machine superiority are inadequate, what sort of guidelines for cognitive task allocation might take their place? We propose a set of principles in which the balance of initiative between human and computer (rather than a static task allocation) is the key concept, and in which advisory prompting plays a crucial role.

In essence, the Advisory function serves as an executive for modulating the balance of initiative between user and computer in the context of an allocation scheme that remains dynamic, flexible, and ultimately under the user's control. The basic mode of operation of the Advisory function depends on some fairly broad characteristics of the task. In particular, there are two main cases:

(1) User Initiative. Under conditions of relatively low workload, at high levels of an organization, and in relatively unstructured tasks (i.e., decision problems that are non-repeating and involve ill-defined options,

outcomes, or goals), the most appropriate allocation mode involves predominant human initiative. The computer's role is to facilitate user-preferred methods of problem-solving. At the same time, however, the computer can back the user up by channeling and advisory prompting. In the latter case, the computer works at least part of the problem in parallel with the user, monitors user behavior (e.g., his use of Select, Modify, Generate and Analyze functions), draws inferences regarding user decision strategies, and provides an Advisory prompt when those strategies seem likely to be significantly suboptimal.

The Advisory function itself is personalized: (A) It does not require users to abandon their natural modes of problem-solving. Rather, it recommends actions that resemble or mesh with user-preferred procedures.

(B) Advisory prompts only occur when the difference between a user-preferred strategy and the solution regarded as optimal by the system is large enough to matter; and the user himself can determine the frequency with which he receives advice, by determining the size of the discrepancy that would set off a prescriptive prompt.

An example of Advisory prompting in this context is the command-level time of fire decision onboard a nuclear attack submarine (Cohen et al., 1982).

(2) Computer Initiative. Under high workload conditions and in relatively structured tasks, the most appropriate allocation mode involves predominant computer initiative; i.e., the computer does the work unless explicitly overridden by the user. The Advisory function in this case requires the computer to monitor its own problem-solving activity (rather than the user's), to assess shortcomings (e.g., incomplete data or conflicting lines of reasoning), and to prompt the user when his contributions are likely to be significant. This form of task allocation exploits a complementarity of expertise in which the human handles unanticipated cues (Johnson, 1985) and knowledge not incorporated in the system (Cohen, 1982),

while the computer simplifies the user's task by focusing his attention on relevant subsets of the problem.

Here, again, personalization is important: (A) Explanations of computer results and requests for user inputs must match user forms of knowledge representation; (B) The user controls the degree of his participation by determining the size of the problem that triggers a call for help.

This form of advisory prompting has been experimentally investigated in the context of air defense target identification (Chinnis et al., 1984).

Nothing prevents both modes of task allocation from applying in different portions of the same decision aid. The route planning application to be discussed in Section 3.0 illustrates such a mixing.

2.6.3 The Appropriateness of Channeling and Prompting. To some extent, a theoretical basis can be provided for the distinction between channeling and prompting. An important factor in designing the prescriptive components of an aid is the nature of the bias to be corrected: in particular, is it the result of an optional strategy, or is it instead a fixed structural feature of the cognitive apparatus. There is disagreement among psychologists on this question. Payne (1982), Beach and Mitchell (1978), and others argue for optionality. They claim that certain natural strategies which appear "suboptimal" are in fact due to an implicit balancing of the costs of errors against the costs (in time and processing capacity) of a more prescriptively correct method. Tversky and Kahneman (1981), on the other hand, have compared cognitive biases to perceptual illusions, which are not influenced by incentives or by conscious awareness. They distinguish two stages in choice: an initial phase of "editing" in which representations are developed for options, outcomes, and their interrelations, and a second stage of evaluation and option selection. Kahneman and Tversky regard the second phase as a set of mostly fixed procedures, and the editing phase, which supplies inputs for the evaluation stage, as the primary source of biases.

The distinction between channeling and prompting corresponds in part to this supposed difference in the origin of a bias. Channeling fosters more adequate problem representations which may either lead to better procedures or improve the results of old procedures; prompting directly encourages the adoption of new procedures. Principles of decision aid design might be developed, therefore, which determine the appropriate mix of channeling and prompting as a function of the degree to which: (a) a bias results from the user's way of viewing the problem; or (b) a bias is the result of optional, controllable processes. We would argue, however, that these are not mutually exclusive. Channeling and prompting might each have potential relevance, if a representation can increase the chance that users will adopt an optional procedure.

In sum, prompting seems appropriate when: (1) errors have potentially significant consequences; (2) the bias is due to optional processes; (3) the recommended procedures resemble and mesh with user-preferred strategies; and (4) the user has requested the advisory prompting function. In short, advisory prompting supplements a user's own implicit cost-benefit analysis, by notifying him when the cost of an error may be higher than he expects; and it reduces the cost of adopting a more prescriptively correct strategy, by encouraging procedures that fit naturally with the preferred approach.

Channeling, on the other hand, seems appropriate when: (1) errors have potentially significant consequences; (2) the bias is caused, at least in part, by a way of representing the problem; and (3) the alternative representation is reasonably natural, even if not spontaneously adopted. Channeling may improve performance by inducing better inputs for a given decision strategy, or by increasing the chance that users will adopt different, more optimal strategies. A final, informal guideline is that since prompting is a more intrusive function than channeling, the threshold for apply-

ing it should be higher, in terms of the significance of errors and user consent.

2.6.4 Alternative Prescriptive Concepts. Channeling and advisory prompting involve the use of prescriptive mechanisms to influence joint user-computer problem solving. However, a variety of alternative prescriptive frameworks now exist for representing and reasoning about uncertainty. A method for handling uncertainty which appears fallacious when viewed in terms of one framework may appear valid (or at least plausible) in terms of another. As a result, these frameworks have different implications for the design of personalized aids, and in some cases, the appropriate forms and occasions for channeling and prompting.

Among the most prominent theories under current debate are variants of Bayesian probability theory, Dempster-Shafer belief functions (Shafer, 1976), and fuzzy set or possibility theory (Zadeh, 1965, 1975). Nonnumerical approaches to reasoning with incomplete information have also received attention, under the general heading of non-monotonic logic (Doyle, 1979; McDermott and Doyle, 1980; Reiter, 1980; McCarthy, 1980) and more recently, the theory of endorsements by Paul Cohen (1983).

These and other theories have recently been reviewed (Cohen, Schum, Freeling and Chinnis, 1984; Cohen, Watson, and Barrett, 1985). For present purposes, we note the differences in the concepts of uncertainty which they attempt to capture; in particular,

- chance or uncertainty about the facts;
- incompleteness or quality of evidence;
- imprecision or vagueness.

Bayesian probability theory, of course, focuses on the concept of chance. DeFinetti (1937, 1964) and Savage (1954) developed formal systems for the

quantification of an individual's "degree of belief," or subjective probability, about uncertain propositions. Von Neumann and Morgenstern (1947) formalized the notion of a subjective dimension of value, i.e., utility. The resulting prescriptive theory has both wide applicability (e.g., it applies both to unique events and to frequency data), and a strong normative appeal: unless beliefs are probabilistically coherent (conform to axioms of probability), one's preferences could conceivably lead to a "Dutch book," in which one loses regardless of the outcome of an uncertain state of affairs. Where the decision-analytic paradigm falls short is in its failure to handle concepts of completeness of evidence and imprecision.

Can a user rely on Bayesian probabilities to tell him how "good" an estimate or conclusion is? Probabilities are not appropriate measures of the quality or credibility of an aid's output. For example, an intelligence estimate that there is a 90% chance that an enemy installation is an SA-6 site is not necessarily better supported than an estimate that puts the chance at 50%. One would not be very comfortable with a conclusion (no matter how high the Bayesian probability) if potentially significant additional data have not been examined. Conversely, the 50% probability could reflect the outcome of a thorough sifting of (possibly conflicting) clues. The acceptability of the conclusion depends on the completeness with which relevant evidence has been consulted, not on the probabilities assigned to the events in question.

In terms of a MYCIN-like rule-based expert system, the chance of a hypothesis is the result of combining "certainty factors" from rules whose antecedents have already been satisfied by data. Incompleteness of evidence, on the other hand, reflects the "shiftability" of that result as more data come in and additional rules are triggered.

When a user describes the altitude of an aircraft as "high," however, this does not necessarily connote either uncertainty or incompleteness of

evidence with regard to its true height. It may be that only a relatively small number of discriminations are required, or natural, for the job at hand. Imprecision is the vagueness with which events, or probabilities of events, are described (e.g., by expert system rules). For similar reasons, probabilities might be described as "very low" or "about 30%." Bayesian probability theory, and the uncertainty calculi employed in expert systems, have been criticized both because they require arbitrarily precise assessments of certainty factors or probabilities, and because the events to which they apply must (in theory at least) be precisely defined (cf., Zadeh, 1983).

2.6.4.1 Implications for Channeling and Prompting. Uncertainty concepts, and the associated normative theories, play a simultaneous descriptive and prescriptive role in the design of aids. On the one hand, acknowledgment of more than one concept increases personalization; it enhances the aid's ability to capture distinctions and insights that are natural to users. Some of the biases imputed to ordinary processes of reasoning appear in a different light when construed in terms of alternative concepts of uncertainty. In particular, a narrow insistence on interpretations in terms of chance may overlook important components of ordinary reasoning that seem in fact to be more concerned with completeness of evidence.

At the same time, however, uncertainty calculi serve as normatively justified constraints, the purpose of which is to improve task performance. We argue, however, that such constraints should be applied through techniques of channeling and advisory prompting that build on, rather than abolish, pre-existing approaches.

2.6.4.2 Incompleteness of Evidence. Several recent prescriptive theories have attempted to explicate the concept of completeness of evidence. Two features of such theories have an important application in understanding natural processes of reasoning:

- Attention is shifted from the statistical likelihood of the hypothesis to the existence of a causal or logical *link* between the data and the hypothesis (Shafer, 1976; Hallden, 1973; Gardenfors et al., 1983; L.J. Cohen, 1977).
- Completeness of evidence depends on the degree of match between the reasons for a hypothesis and a domain-specific list of potential reasons (L.J. Cohen, 1977; Paul Cohen, 1983).

Research on human processes of reasoning about uncertainty (Section 2.5 above) suggests that each of these factors plays a role. For example, the "conjunction effect," in which detailed hypotheses are judged more probable than their components, may involve the application by users of knowledge structures which define what constitutes a satisfactory causal explanation within the domain in question (Leddo et al., 1984). A second example is in the effect of outcome values on probability assessments. According to Einhorn and Hogarth (1984), if an outcome is positive, its probability is likely to be underestimated; if negative, overestimated. The extent of the bias is dependent on the incompleteness of evidence for the probability in question.

In these examples, users may have reasonably valid assessments of the completeness of evidence for a hypothesis, but they may have severely biased assessments of the chance of its being true. The validity of natural reasoning, and the potential user contribution to collaborative problem solving, depends, therefore, on which interpretation the computer adopts. A major function of channeling and prompting should be to keep relevant concepts of uncertainty distinct, and to focus user judgments on the applications to which they are appropriate.

Some tentative conclusions with regard to the design of personalized aids are as follows:

• Statistical frequency-based inference should be handled by the computer. Users may contribute to reasoning based on causal mechanisms (cf., Johnson, 1985).

- Explanations and recommendations offered by the aid should be formulated in causal terms whenever possible, and should match user knowledge structures which are naturally associated with inferences about the events in question. Such user knowledge structures may be thought of as frames with slots corresponding to the types of evidence regarded as relevant to the hypothesis in question. When the evidence for a probability estimate is "representative" (Kahneman and Tversky, 1972) of the expected kinds of evidence, user confidence in aid outputs should be increased. Reliance by the aid on atypical bits of evidence will detract from confidence in the assessment (Leddo et al., 1984).
- Displays and prompts should clearly distinguish between issues related to chance and issues related to completeness of evidence. For example, the system may provide estimates of the expected utility of an option, but at the same time indicate what new information, should it be obtained, might force a reevaluation or revision of the option. The system may display probabilities (i.e., chances) for different possible classifications of a SAM site, but prompt users when the evidence for those probabilities is incomplete, or when conflict among different bits of evidence or lines of reasoning casts doubt on the reliability of the probabilistic estimates. In these cases, prompts may also notify users when users themselves are likely to have a significant contribution -- by incorporating contingency plans within an option, by providing additional evidence with respect to a hypothesis, or by revising some of the assumptions and beliefs that led to a conflict. Such prompts might simplify the user's task by suggesting candidate plans or revisions for selection by the user (Cohen, Watson, and Barrett, 1985).
- 2.6.4.3 Imprecision. The use of fuzzy variables as inputs and/or outputs for a decision aid might significantly reduce the assessment burden usually associated with prescriptive models. Watson et al. (1979) have described the use of fuzzy linguistic variables (e.g., "about 30%") as surrogates for probabilities. Techniques of this sort might be applied to modifications by users of database values, user-provided constraints on option generation, and user-specified queries or requests for information.

A new line of research, with an important role in channeling, might involve fuzzy relations or structures. Typically, expert systems and decision-analytic aids ignore nterdependencies among data and hypotheses, due to

the number and complexity of the assessments that would otherwise be involved. Such interdependencies, however, are potentially critical in many inferential problems. It may be possible, however, to represent complex interdependencies quite simply in terms of fuzzy relations (e.g., "highly corroborative," "slightly redundant"). A simplification of this sort might facilitate user attention to prescriptively critical factors that would otherwise be too complex to handle.

When relatively large numbers of user inputs are required by an aid, Advisory prompts might work in tandem with fuzzy representations to tailor the input dialogue to a user's preferred style. Thus, the aid might extract the maximal conclusions possible from a user's initial imprecise inputs. If these conclusions are insufficiently precise to establish a desired conclusion, Advisory prompts might guide the user through the minimal set of judgments required to extract an answer that has the required degree of precision. Such judgments, to the extent possible, would be merely ordinal (i.e., which event is more probable?) and/or fuzzy (e.g., is A "very much more likely" than B?).

3.0 ROUTE PLANNING

The feasibility of developing a personalized decision aid, within the framework of the last chapter, has been explored in the context of tactical air route planning. In this chapter, the initial steps are taken--by using that framework to characterize the route planning decision task as currently performed by the Wing Operations Center (Section 3.1); and by using that framework to examine (Section 3.2) and critically analyze (Section 3.3) a prototype decision aid for route planning developed by SCT. The results of this discussion are the starting point for a new design concept to be presented in Section 4.0.

3.1 The Decision Task

In route planning, wing or squadron mission planners select the routes that will be followed by aircraft from friendly territory to a target and back to friendly territory. Participants may include actual pilots of the aircraft to be flown, intelligence officers, and aides. The present work focuses on offensive counter-air and interdiction missions for F-111 aircraft. The time available for route-planning in this (and other) cases is typically highly constrained: nominally it may be as long as 2 hours; in fact, it may be as brief as 15 minutes.

The components of the route-planning task match the phases of problemsolving outlined in Section 2.5:

Identification of Objectives: The planning process is initiated by an Air Tasking Order which specifies the target, time over target, ordnance, and other details (e.g., call signs, refueling points, etc.). Additional constraints are included in Air Space Control Orders: i.e., safe passage corridors or control points which govern transit over friendly territory.

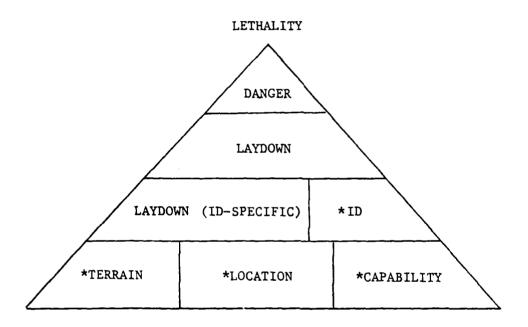
Situation Assessment: Intelligence officers provide information on threat locations and identities. Many of these assessments are highly uncertaindue to threat mobility (perhaps 20-30% are non-fixed), shifting enemy usage, possible recent destruction, and datedness of reconnaissance and intelligence reports. In the route-planning session, pilots may contribute information not yet available to intelligence, based on their own recent observations.

Option Generation: A route is plotted based on information about threats and terrain contributed by pilots and intelligence officers. Typically, the route is generated in stages: first, an initial point (IP) near the target is selected to serve as a navigation update prior to the final straight path run at the target, called the attack axis. Then an ingress route from the FLOT (Forward Line of Own Troops) to the IP is plotted. Next the egress route, from the target to the FLOT, is determined. In the final stage, approach and return routes over friendly territory are specified. Other aspects of the mission, such as speed and altitude, weapons configuration, and potential use of countermeasures (such as jamming and chaff) may also be identified.

Choice: Refinements in candidate routes (or portions of routes) may involve considerations such as leeway for fuel, better use of terrain or landmarks, revised threat assessments, potential safe areas (if the mission were aborted over hostile territory), and secondary targets. All of these matters involve a balancing of multiple objectives: viz., (1) to destroy the target and (2) to get back safely; and they may elicit considerable give-and-take among intelligence officers and pilots.

Route planning, in sum, is a relatively demanding decision-making process, performed in a group setting by diverse individuals under rather severe constraints of time. It involves extreme uncertainty, complex options, and multiple objectives. In all these aspects, it seems a promising candidate for personalized decision aiding.

Figure 4: A hierarchy of uncertainty riables



* = Inputs to the route planning system

uncertainties at that level. TOPIC and FOCUS operate jointly to determine map displays.

The inferential process underlying route planning is schematically depicted in the "inference pyramid" of Figure 4. It begins with basic data regarding terrain, threat locations, threat identifications, and threat capabilities, and culminates in overall measures of lethality for routes. "Basic data" may in fact be the outputs of inferential processes in other aids (e.g., for identifying and locating enemy C³ installations and tracking enemy order of battle) and will themselves involve uncertainty.

Under TOPIC, the user may select information from any "block" in this pyramid:

- TERRAIN: Terrain contour maps are displayed for an area determined by the current FOCUS (e.g., OVERVIEW (Figure A-8), REGION, a THREAT vicinity, etc.).
- LOCATION: Ellipses are displayed indicating the region within which a threat is thought to be located with 95% probability (Figures A-9, A-10).
- CAPABILITY: The range and envelope of a threat are indicated as a contour which shows the distance at which an aircraft would incur a 50% chance of being destroyed by that threat. When the ID of a selected threat is uncertain, capability contours can be displayed for each of the specific threat-types which are considered possible classifications of the threat (Figures A-11, A-12). A contour is also available which represents the overall expected capability of the threat, based on the probabilistically combined capabilities of the different possible threat types (Figure A-13). CAPABILITY is a somewhat abstract assessment: it ignores terrain and location uncertainty.
- ID: Probabilities are displayed, in the form of a histogram, for each possible classification of a threat (Figure A-14). The probability for "None" represents the chance that no threat exists in the relevant area (or that what appear to be two threats should in fact be regarded as one).
- LAYDOWN-ID: This display represents the danger of a threat, assuming its identity to be of a specified type; it combines information about the terrain in the vicinity of the threat, informa-

- i.e., if a selected route crosses a region, then when that REGION is selected as a FOCUS, the relevant portion of the route will be shown (Figure A-6); similarly, in ANALYZE displays, evaluation data will appear for selected routes only.
- SEGMENT: The user can designate any portion of a route for close-up display by pointing with the cursor to the beginning of the segment and selecting; then pointing to the end of the segment and selecting again (followed by "escape"). Alternatively, if the user selects a threat, the aid displays the route segments whose design was determined, in part, by the presence of the designated threat. These displays provide useful information on the selected segment and the region through which it passes: e.g., possible threat identities, masked regions, headings and location data, and--where relevant--lethality envelopes showing areas where a small amount of divergence from a route is expected to increase lethality by 20% or more (Figure A-7).
- LEG: Routes are automatically segmented by leg (i.e., portions of a route between maneuvers); when he points at a leg, the user creates a display centered on that leg.
- GRID: Selecting GRID results in the display of a longitude and latitude grid on all relevant displays; selecting GRID again removes them; and so on. Grids can be switched on or off without otherwise changing a selected display.

Users have maximal flexibility, through FOCUS, in the design of displays: any region, route, leg, route segment, etc., can be shown. At the same time, however, FOCUS facilitates the organization of knowledge according to prescriptively meaningful categories: i.e., THREAT produces a "zoom" on just that region where a particular threat is effective; SEGMENT can focus in on just those portions of routes that are the way they are because of a particular threat.

4.6 TOPIC

TOPIC determines the specific nature of the information to be displayed about a selected object. In particular, it enables users to tap into an inferential process at any level and observe the conclusions and

- a mix of automatically provided constraints (from the ATO or ASCO) and manual inputs or revisions is facilitated;
- use of graphical input methods is maximized;
- the CONSTRAINTS module is not regarded as a once-and-for-all "Initialization" process; users may return to it, make changes in mission parameters, and observe the effects of the changes on the evaluation of a route. In effect, then, it facilitates regarding these parameter settings as parts of a single complex decision that must be made in the course of mission planning.

4.5 FOCUS

In FOCUS, users can select an "object" about which information is to be displayed. Potential objects include overview (i.e., the entire region designated under CONSTRAINTS) or a specified region, threat, route, route segment, or leg. Choices are made graphically. When FOCUS is selected, two things happen: (1) an overview "menu" display appears which shows all threats (by number) and all previously generated routes; and (2) a submenu appears listing the types of object available. The user makes a selection from the submenu, then from the overview screen, by one of the following methods:

- REGION: A rubber-band box appears on the overview screen (Figure A-2); its location and size can be adjusted by moving the mouse (sizing occurs when the left function key is depressed). A region is actually selected when the user depresses the "escape" (right) function key. The display then "zooms" in on the chosen area, providing details not available in the overview: e.g., possible threat identities and areas masked from a threat by terrain (Figure A-3).
- THREAT: An arrow cursor appears, with which the user selects the desired threat. A new screen then "zooms" in on the selected threat (Figure A-4). Again, detailed information is provided.
- ROUTE: The user points at a route on the overview menu and selects it with the left cursor key (Figure A-5). Any number of routes may be selected in this way. Routes may be unselected by pointing and depressing the middle ("erase") cursor key. The process is terminated by pressing the "escape" (right) key. Selected routes are displayed thereafter on all relevant screens:

9011	•
REPORTS	
ANAL YZE	LETHALITY EXPLAIN FUEL ALERTS ISEGMENT ITHREAT ELIMINATE
GENERATE	ATTACK INGRESS EGRESS POINTS DRAW AUTO ALT/SPD CONTINGENCY
MOD I FY	CPARAMETER A> CPARAMETER B> CNDG UNDG DEFAULT EXPLAIN SENSITIVITY MEXT PRIOR
10p	DANGER LAYDOWN LAYDOWN·ID ID TERRAIN LOCATION CAPABILITY THREAT·TYPE
Focus	OVERVIEU REGION THREAT ROUTE SEGHENT LEG GRID
CORSTRAIRTS	

Figure 3: Menu structure of the Personalized Route Planner.

time through CONSTRAINTS; and the result will be reflected in subsequent displays under all other functions.

Figure 3 gives the second-level menus for each of the top-level items. They will be discussed in turn in the following sections. Sample display screens for each of these modules are contained in the Appendix, and will be referenced specifically at appropriate points in the text.

4.4 <u>CONSTRAINTS</u>

In this module, users may input, view, and revise mission constraints and objectives (Figure A-1). In particular, they may:

- o type in display coordinates for the region of interest (at the upper left of the display);
- o examine mission parameters specified by the Air Tasking Order and Air Space Control Orders (data on right-hand side of the screen) and electronically transmitted to the route planning aid;
- o manually input or revise these mission parameters by typing in the blank spaces; user inputs (e.g., target control points over friendly territory) will then be automatically displayed on the map. An alternative, graphical input method for some parameters would involve selecting "Target" (with the mouse cursor and left mouse function key), then selecting the desired location on the map; the appropriate longitude and latitude would then be displayed automatically in the blanks on the right.
- o manually input mission parameters which are determined at the wing or squadron level (e.g., aircraft to be used, its weight under the selected configuration, availability of jamming, visibility and climate conditions), by selecting appropriate values with the mouse.

Although the details of this display are highly provisional, several points are worth noting:

o the "form-filling" method is more flexible than a linear, question-and-answer dialogue: inputs can be provided, or revised, at any time and in any order;

4.3 Overview of Menu Organization

The first menu which a PRP user sees has the following items:

CONSTRAINTS FOCUS TOPIC MODIFY GENERATE ANALYZE REPORTS QUIT

Most of these items, when selected, produce a second-level menu of actual operations to be performed on the display (by the computer or by the user). Those operations will be the subject of the following sections. Here we comment briefly on the overall structure and operation of the top-level menu.

It will be noted that these top-level choices roughly correspond to the "stages" of problem solving outlined in Section 2.5. To this extent, they represent a natural high-level segmentation of the route planning task. Thus, CONSTRAINTS enables users to view or input the objectives of the mission, as specified in the Air Tasking Order; FOCUS and TOPIC request the display of objects and information about those objects, respectively, in the situation assessment process. (MODIFY allows users to adjust that information.) GENERATE permits user and/or computer specification of route options. ANALYZE involves the evaluation and selection of an option for execution. (REPORTS provides for the production of hard-copy outputs based on information generated under the other functions.)

While the top-level menu has a natural organization, however, it is by no means a rigid script. Users may skip freely from one section of this structure to another. For example, TOPIC and FOCUS displays can be used at any time, (a) to view information, (b) to MODIFY information, or (c) to GENERATE a route against the backdrop of the selected TOPIC/FOCUS display. Once a route has been generated, (a) it can be evaluated using special charts in ANALYZE, or (b) it can be displayed against a map-like TOPIC/FOCUS display. Finally, mission parameters can be changed at any

- The mouse affords direct linear spatial control over a cursor by the corresponding spatial movements. It is a highly efficient and natural pointing device (Card, Moran, and Newell, 1983).
- The mouse lends itself well to a variety of interactive graphics functions: e.g., "rubber banding" in drawing lines, "rubber boxes" for indicating regions, etc. (Foley, 1981).

The three mouse keys are used consistently throughout PRP to represent three basic functions: selecting, erasing, and escaping. The left key is used for selecting options or drawing, the middle key for deleting, and the right key for escaping to the next higher logical level, i.e., from the graphics screen to the menu, or from a menu to the next higher menu.

Dialogue design. In PRP, the user may go anywhere in the system, at any time, from anywhere, quickly and without losing relevant selections or displays. There is a very shallow menu structure (typically only two levels), and users may "pop up" quickly to the top menu by use of the "escape" mouse key. This creates the effect of a "palette" of information-processing operations all simultaneously available to the user with respect to a given task.

A menu line is always present on the display, together with a help line which describes the menu function to which the cursor is pointing. Note the advantages of the menu format over other forms of interaction:

- it is fast for experienced users (unlike computer-initiated question-and-answer dialogues); and
- it provides support regarding available options for inexperienced users (unlike user-initiated command languages).

Error correction is facilitated in a variety of contexts: e.g., routes or route segments can be drawn and subsequently erased; adjustments to parameters can be undone; and default parameters can be restored after any number of adjustments.

Also, we have not included cultural details (towns, landmarks, etc.) which would enhance the value of map displays.

PRP is implemented on an IBM-PC/AT with 512 K-bytes of RAM and a 20 mbyte hard disk. Other features of the system are: an IBM Color Graphics Adapter in 640x200 two-color resolution, a Princeton Graphic Systems HX-12 monitor, a Mouse Systems optical mouse, and an IBM communications adapter. The demonstration system utilizes a single display screen. We are convinced that this is sufficient, in conjunction with the present PRP design, to support the route planning process effectively. However, a variant of the present design could easily be extended to a system incorporating two graphics monitors. Although the present system is implemented in black and white (in order to maintain high resolution), we would anticipate use of color in a final design; e.g., for such purposes as distinguishing danger contours at different levels of risk.

4.2 Overview of the User Interface

The interface features of the Personalized Route Planner are intended to minimize the attention users must devote simply to operating the aid. Here we outline some of the features that will appear in more detail in later sections.

Displays. Virtually all displays present information graphically, by a combination of maps and charts. Wherever possible, correlations between different displays involve natural spatial or symbolic correspondences.

Input modes. All user inputs are by means of a single input device: a mouse, with three function keys. (A keyboard is, however, optional for inputting data from the Air Tasking Order.) The advantages of this arrangement are:

• Time-consuming transitions among input devices are eliminated.

4.0 A PERSONALIZED ROUTE PLANNING AID

A demonstration system has been implemented in the route planning context to illustrate, and to demonstrate the feasibility of, personalized decision aiding. The demonstration system encourages a highly flexible, personalized style of problem solving, while channeling and prompting users toward more prescriptively optimal strategies. In so doing, it addresses most of the shortcomings raised in the previous chapter in regard to the present version of the Route Planning Aid.

4.1 <u>Implementation</u>

The Personalized Route Planner (PRP) is by no means intended to represent an operating decision aid. It consists of about 65 canned screens embedded within a "live" menu system. The screens provide appropriate displays for a large number of menu requests, representing a specific route planning example. Input and output operations for some displays are also operational. A user who stays within the broad boundaries of the example may, therefore, get a fairly good feeling for the intended operation of the aid.

A demonstration system of this sort can be quite useful as a tool for design, research, and evaluation. It represents a form of limited prototyping, in which important features of a top-level design are simulated; they may then be subjected to scrutiny by potential users and (if necessary) redesigned. These features include the basic functions of the aid, the allocation of tasks between users and the system, and user interface functions. Computational mechanisms are not implemented, although thought has been given to their eventual design.

The displays in the present implementation are not intended to be definitive in their details. For example, we have omitted indicators of latitude and longitude which would appear in a final, higher-resolution version.

components clutter the graphics screen and are easily confused with numbers for threat sites.

- <u>Input modes</u>: RPA requires three separate response devices: a trackball and two keyboards. Transitions among these devices may significantly slow the interactive process.
- <u>Dialogue design</u>: RPA has a relatively inflexible, compartmentalized structure. For example, threat information is no longer available during route generation. Error correction of any kind is difficult. A rigid sequence is imposed on all operations-initialization, threat briefing, attack axis generation, ingress route generation, egress route generation (though iteration is permitted at a few specific points). It is not possible to change scale or zoom in on a region of interest during an analysis.

Speed will be a major factor in evaluating any aid for route planning. In an empirical evaluation, route planning was found to be slower with RPA than without it (Adelman and Crowley, 1984). The primary cause of this poor performance may have been the slow processing speed of the aid's optimization algorithms and/or rule base. However, we suspect that some basic human engineering considerations, such as those raised above, also have played a role.

may make sense if only the ingress is considered, but it may lead to a shortfall in fuel on the egress.

- o RPA does not permit simultaneous specification and evaluation of speed, altitude, weapons configuration, and route. A decision on one of these has repercussions for the evaluation of the others.
- o RPA's route generation algorithm relies exclusively on terrain and threat capability data. "Minimum lethality" is achieved, therefore, only in a very narrow sense. Among the ignored factors are: (1) fuel constraints; (2) required lead time into a turn (legs that are too short); (3) predictability (legs that are too long); and (4) avoiding maneuvers in heavily defended areas. In the evaluations of RPA, users did not feel comfortable having to incorporate these factors on their own.
- o RPA displays only a single "optimal" route. It does not support user revisions by showing route "corridors," or potentially acceptable alternative routes.

Choice.

- o RPA provides a single evaluative measure for a route: lethality. It thus fails to support the balancing of two route-planning objectives: survival and destruction of the target. In essence, what RPA misses is that there is a higher premium for avoiding risk on the ingress (before reaching the target) than on the egress.
- o In the absence of a clear interpretation for RPA's numerical outputs, the user has no way of determining how much a given "lethality" difference between routes matters. There is also no clear relationship between the "lethality" of an ingress route, the "danger index" for a leg, and the "worry" indicator for a threat.

Interface: A distinct set of problems for RPA pertains to the user interface. These criticisms fall roughly into three categories:

o <u>Displays</u>: Correlating information from RPA's two screens is difficult. There is a natural mapping between the graphical displays used to show threat laydowns, danger contours, and terrain contours, and the textual tables used to provide route explanations and critiques. Numbers used to identify route

planning context. At the same time, it has shortcomings which, we would argue, may strongly affect its ultimate usefulness. Some of these problems are mentioned in an evaluation of RPA by experts and non-experts (Gates and Figgins, 1984); others emerge from an application of the concepts developed in Section 2.0 above. Our purpose in discussing them now is not so much to criticize RPA, as to introduce some positive ideas, based on the above framework, for a revised design. These points will be elaborated in the discussion of a Personalized Route Planner Aid in Section 4.0.

Situation assessment:

- o RPA fails to acknowledge or deal with critical uncertainties regarding threat locations, types, and capabilities. In this regard, as well as others, it may replicate the unaided route-planning process too closely. Decisions which take uncertainty into account may differ dramatically from decisions which are based on a "best guess."
- o The numbers produced by RPA as measures of danger or lethality have no clear interpretation. As a result: (1) it is difficult for users to second-guess or override aid outputs; (2) a consistent numerical framework does not exist for exchanging inputs and outputs with other decision-aiding systems; and (3) the validity of RPA's outputs (even in a comparative sense) is subject to doubt.
- o The aid does not permit users to override or adjust default values. On-the-spot user contributions, e.g., information from recently returned pilots regarding the location of a mobile threat, may be extremely valuable.

Option Generation.

o RPA forces routes to be generated and evaluated piecewise. The attack axis, ingress route, and egress route are generated separately; no lethality measures, explanations, or critiques are presented for the route as a whole. Iteration is possible, but only at specific points. Although it replicates unaided practice, RPA thus fails to support the handling of interactions among route components: e.g., by-passing a threat on the ingress

3.2 The Current Route Planning Aid

A starting point for the present work is the prototype Route Planning Aid (RPA), developed for RADC by Systems Control Technology, Inc. (SCT). RPA is described in a variety of reports (Riemenschneider and Rockmore, 1983; Riemenschneider, Rockmore, and Wikman, 1983; SCT draft, 1983), and will not be described in detail here. For present purposes, however, we note the following features:

- <u>Situation assessment</u>: RPA contains a threat database, which includes the location (latitude/longitude) and type of each threat; it utilizes a lethality model for each threat type, which specifies danger to an aircraft from that threat as a function of cross-range, down-range, speed, and direction. These data are used to generate a "threat laydown" display, i.e., a map of threat locations showing range envelopes around each threat; and a "lethality contour" display, which shows danger levels in the mission area. A terrain contour display is also available.
- Option generation: Based on this threat information and on a user-specified target, friendly airspace exit point, and friendly airspace reentry point, RPA employs an optimization routine to determine "minimum lethality routes" from the ingress point to the target and from the target to the egress point. Option generation takes place in stages: first, attack axis, then ingress route, then egress route. At each phase, users may generate their own route to accommodate factors omitted by RPA's lethality model.
- Choice: In generating their own routes, users may be assisted by a numerical lethality measure for the relevant route component, a textual explanation of the threats that affect that component, a more detailed leg-by-leg description of threats, and by a route component critique. The latter is produced by a heuristic rule base within RPA, and describes factors not taken into account by RPA's automatic route generation process. Based on the automatic route and on this critique, users are expected to manually generate routes that accommodate the additional factors while not exceeding the lethality of the automatic route by "too much."

3.3 Problems with RPA

RPA includes a number of promising ideas for decision aiding in the route

tion on its probable location, and data about threat capability. Danger is represented as a set of contours showing the distances at which an aircraft would face a 20%, 50%, or 80% chance of destruction (Figures A-15, A-16). The distinction between LAYDOWN-ID and CAPABILITY is that LAYDOWN-ID, unlike CAPABILITY, incorporates geography: LAYDOWN-ID takes account of location uncertainty regarding the threat and terrain contours in its vicinity.

- LAYDOWN: LAYDOWN-ID is probabilistically combined with information about the likelihoods of different possible threat IDs, to determine the overall expected danger of the threat. This is a set of 20%/50%/80% danger contours which factor in terrain, location uncertainty, capabilities, and ID uncertainty (Figure A-17). For the overview display, only 50% contours are shown (Figure A-18).
- DANGER: While LAYDOWN treats each threat separately, DANGER contours represent the probabilistic integration of LAYDOWN contours for different threats. They reflect the risk to an aircraft of being in a certain location from all threats in that vicinity (Figures A-2, A-4). Since danger incorporates information from all levels below it in the pyramid, it is the default setting in TOPIC.

At the very top of the pyramid, lethality (which appears in the ANALYZE module) is the cumulative danger of a route, obtained (in part) by integrating the DANGER measures along the locations to be traversed by an aircraft on that route.

THREAT-TYPE is a general threat database. It allows users to examine capability contours for any threat-type without regard to possible classifications of a particular threat.

The TOPIC displays highlight the existence of uncertainty about threat identities, locations, and capabilities; and facilitate an awareness of the inferential steps that are required for any particular conclusion in the route planning process. As the pyramid is ascended, an increasingly comprehensive set of factors is accounted for. In addition, the breakdown of CAPABILITY and LAYDOWN-ID according to potential classifications of a

particular threat facilitates the exploration by users of the implications of ID uncertainty.

4.7 MODIFY

MODIFY enables user to quickly and easily insert their own judgments into the situation assessment process, if they choose to do so. Adjustments in conclusions may be made at any level of the "inference pyramid." They may be based on on-the-spot evidence not yet incorporated in the computerized analysis (e.g., direct observations by pilots of mobile threats), or they may be used for what-if analysis. When a modification is made, all computations above it in the pyramid are affected accordingly; values below it are still computed, but do not flow upward.

The adjustment process is based on techniques of interactive graphics and a spatial representation of each parameter to be adjusted. MODIFY opplies to whatever TOPIC/FOCUS display is currently available. The MODIFY procedure varies somewhat, however, for different variables under TOPIC, depending on the most natural decomposition of that variable into spatially represented parameters. The first 1 to 3 choices in the MODIFY submenu (Figure 3) reflect these specialized procedures, while the last six choices are common to all TOPICs. We first consider the specialized methods for modifying variables:

• MODIFY (CAPABILITY): Threat capability is represented as a pie, with three parameters: RANGE (i.e., radius), SPREAD (i.e., the angle corresponding to the width of the threat envelope), and ORIENTATION. Selecting any of these three items means that movements of the mouse will cause changes in the corresponding parameter (up/down to increase/decrease RANGE and ANGLE, left/right to change ORIENTATION). In this (as in all MODIFY procedures), the original values continue to be displayed as a reference, while changes are represented symbolically in relation to the original values. (Figure A-19 shows a tentative increase in the range of a threat.) Changes are finalized (subject to a subsequent UNDO) by pressing the left mouse function key.

- MODIFY (LOCATION): Threat location is represented by a 95% uncertainty ellipse, which is broken down into three parameters: POSITION, AXES, and ORIENTATION (the latter is not implemented in the demonstration). After POSITION is selected, mouse movements cause corresponding movements in the ellipse; by selecting AXES, users may alter its eccentricity (by vertical mouse movements) or radius (by horizontal mouse movements). Before finalizing an adjustment, changes are represented by crossed lines (Figure A-20).
- MODIFY (ID): Any or all columns of the probability histogram may be adjusted by moving the cursor to a desired position on the vertical scale above the relevant classification type, and pressing the left function key (Figure A-21). In doing so, users need not be concerned that ID probabilities, represented by the heights of the columns, add to 1. They can focus on the relative likelihoods of different ID possibilities, i.e., the relative heights of the columns. When they are done, selecting CALCULATE causes the aid to normalize the adjusted values automatically.
- MODIFY (LAYDOWN-ID, LAYDOWN): Users may directly modify laydown contours for a threat by drawing and erasing with the cursor. This function can also be used to draw in a new threat or area of danger.
- MODIFY (DANGER): In a similar manner, users may directly revise danger contours.

A variety of other functions support the MODIFY process, and are available for all TOPICs:

- UNDO: Users may restore the values existing prior to the current MODIFY interaction. Those values themselves may be due to previous user adjustment.
- DEFAULT: Users may restore original (prestored or automatically computed) values.
- EXPLAIN: Reasons for the current values are displayed. Computergenerated values trigger an automatic explanation facility (e.g., evidence and rules upon which an assessment is based). Users can employ this function to record and retrieve the reasons for their own adjustments.
- SENSITIVITY: This function enables users to gauge the importance of an adjustment in the route evaluation process. It graphically displays the impact of variations in the selected parameter on

lethality for two user-selected routes (Figure A-22). It shows how much adjustment in the parameter would be required to make a difference in the choice between those routes, by highlighting the current value of the parameter and the crossover point between preference for one route and preference for the other. Since overall lethality and lethality for the ingress portion of the route are separable factors in the evaluation of a route (see Section 4.9 below), they are each plotted.

• NEXT/PRIOR: Users may wish to examine or modify additional regions, threats, legs, or route segments without having to return to FOCUS in order to make a new selection. NEXT moves the user to the next item of the same type in a logical sequence (i.e., the next region of the same size in the OVERVIEW area, the next threat in the vicinity, the next leg along the same route, the next route segment whose design was influenced by a threat). PRIOR moves the user back one step in the same sequence.

4.8 GENERATE

The GENERATE module contains a highly flexible set of functions for identifying potential routes. It offers the options of manual or automatic route specification; more importantly, it permits a blending of user constraints and automatic route generation to any degree from completely manual to completely automatic. In addition, routes may be generated or revised in the traditional order (attack axis, ingress route, egress route, friendly approach and exit), as a whole, or in components arbitrarily selected by the user. Any maplike TOPIC/FOCUS display may be used as the background for a route to be generated.

The first group of items in the GENERATE submenu (ATTACK, INGRESS, EGRESS, POINTS, and DRAW) are all methods by which users constrain routes. Such constraints may be fuzzy and partial, or they may fully define a candidate route. When AUTO is selected, the system produces a set of routes which respects whatever constraints the user has provided. User constraints are specified via interactive graphics techniques and may take the following forms:

- ATTACK: Users select potential initial points with the mouse, producing a symbol on the display for each potential IP (Figure A-23). If only one IP is designated, that IP will be used in the automatic route generation process (AUTO). However, any number of potential IPs can be identified by the user, for consideration by the automatic process. If no IP is specified, one will be selected by AUTO. Users may get a detailed view of the region surrounding the target by use of FOCUS/REGION; such a display may then be used for the selection of IPs. Whenever ATTACK is selected, a shaded disk appears on the display showing the region which is between, say, 20 and 30 miles distant from the target.
- INGRESS/EGRESS: Users specify potential ingress and/or egress locations. These locations may be intervals, rather than points, and any number may be identified. The input method is simply a matter of drawing (and erasing) intervals with the mouse.
- POINTS: Users may indicate a set of points through which, or near which, the route must pass (Figure A-24).
- DRAW: Users can fully specify portions of a route simply by drawing them. Any automatically generated routes will include these segments as a part. Portions of previously generated routes may be used to constrain new routes, by first displaying the old route, using DRAW to erase the undesired portions, then calling AUTO.

AUTO incorporates these constraints into a knowledge-based, heuristic route generation process. AUTO will generate a complete route or a portion of a route, depending on the user inputs which have been provided. Thus, users may first use ATTACK to select an IP. Then, if ingress intervals are specified via INGRESS, AUTO will generate an ingress route. If the user then goes on to specify egress intervals (leaving the previously generated attack axis and ingress route on the screen as constraints for AUTO), a complete route (including the earlier components) is produced. In other words, AUTO facilitates, but does not require, the traditional route generation sequence. On the other hand, simultaneous generation of the entire route is also facilitated: if IPs, ingress intervals, and egress intervals are all specified before selecting AUTO, a total route will be generated. Finally, users may concentrate attention on a small portion of a route by

FOCUSing on any desired region and designating two control POINTs within it; AUTO will generate route segments connecting those points.

In this regard, an Advisory prompting function is valuable. A user who adheres to the traditional order of route generation may generate routes that fall significantly short of optimal. For example, an ingress route which is the safest way to get from the FLOT to the IP may not in fact be the ingress that would be chosen when one takes account of fuel problems arising in the egress. PRP provides a message to users who have generated a route piecewise, when significant improvements would be obtained by generating a route all at once under the same constraints.

GENERATE facilitates consideration of multiple routes: AUTO gives users the choice of seeing 1 to 4 different routes which satisfy the indicated constraints. The heuristic route generation process ensures that these routes are qualitatively different (e.g., adopt different strategies with regard to a particular threat), rather than minor variants. An Advisory Prompt notifies users who have requested only 1 route, when alternatives exist which are significantly superior on at least one such qualitative dimension.

GENERATE facilitates improved decision making in two additional ways: by supporting simultaneous consideration of other aspects of the mission planning decision (altitude and speed), and by prompting for the formulation of contingency plans.

- ALT/SPD: Like route generation, decisions regarding altitude and speed blend user and computer contributions. Thus, users may specify altitude and/or speed for any portion or portions of a route. To do so, they point at the relevant spatially located speed or altitude symbol, use the left key to increase altitude or speed, and use the middle key to decrease it (Figure A-25). AUTO will respect these constraints, but provide its own altitude and speed recommendations for unconstrained portions of the route.
- Prompting for CONTINGENCY Plans: In some cases, uncertainty about a parameter (such as the ID of a threat) is a key factor in the

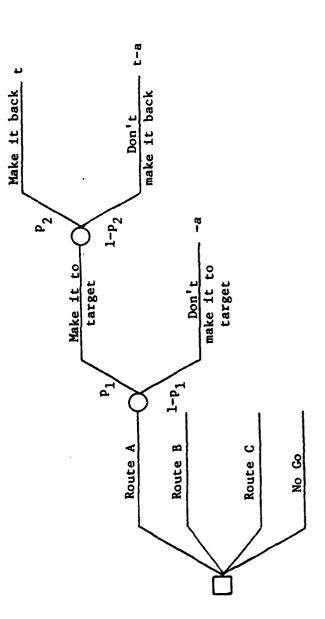
generation of a route. For example, Route A could be considerably shortened if Threat 6 were known to be an SA-9 rather than an SA-10. When uncertainty is critical in this way, PRP examines the evidence for the current assessment of that parameter. If it is relatively complete, no action is taken. On the other hand, customary data bearing on the parameter may be missing: e.g., EM sensors in the vicinity of Threat 6 have been destroyed; no recent overflights near Threat 6 have been conducted. In that case, an Advisory prompt is triggered in conjunction with the generation of Route A (Figure A-26). Such a prompt says, in effect: if there is a chance of obtaining information on the parameter (e.g., Threat 6 ID) during the course of the mission, contingency plans for that possibility should be considered now.

• CONTINGENCY: The user may respond to this prompt by selecting CONTINGENCY. That function displays the choice points along a route where decisions might occur based on new information. For example, if evidence were obtained early in Route A that Threat 6 is an SA-9, a contingency branch of Route A takes the aircraft closer to Threat 6, on a shorter route to the target (Figure A-27). As a result of fuel savings on the ingress, a second contingency is also created in which the aircraft takes a longer route on egress in order to avoid Threat 16.

4.9 ANALYZE

ANALYZE helps users compare routes and make choices. In doing so, it draws on a prescriptive model of route selection. The objective of ANALYZE is to communicate the essential implications of that model, by dramatizing and explaining important differences between routes, without requiring users to grapple with its details.

A simple prescriptive characterization of route evaluation is given by the decision tree in Figure 5. It represents a choice among 3 routes (A, B, and C) and an option of not attacking at all. For each route, there are two important uncertainties: survival up to the target, and survival on the return. (Other uncertainties, such as the portion of the target which will be destroyed, are ignored in this simple model; they can be incorporated easily without changing the essential points.) Each of these uncertainties is characterized by a probability: p_1 = the chance an aircraft



しょうしょう ちゅうしょ しょうりょう 有事的 アイア アイア 日本 日本 アイア アイフィッション アルコン アイアイ オイス オイス アイ・スティア・ア

Value of Route A $= p_1 t - (1-p_1 p_2)a$

Figure 5: A simplified prescriptive model for route planning.

TORNESS OF CONSTRUCT PRODUCT TORNESS TO

will make it to the target; p₂ - the chance an aircraft will make it from the target back to its recovery base. Further, each outcome is associated with a value: t - the value of a successful strike at the target; a - the loss involved in the destruction of the aircraft.

It is illuminating to compute the value, or "expected utility," of a route in this model. In words:

Value of Route A = (the probability of a successful ingress) x (the value of the target) - (the probability of the aircraft being destroyed anywhere on the route) x (the value of the aircraft).

This equation highlights two important features of the prescriptive approach to route planning: (1) it requires a comparison between target value and the value of friendly forces; and (2) it distinguishes between risks on the ingress and risks associated with the entire route. In essence, what this equation says is that the chance of damaging the target (i.e., success on ingress) and the value of the target must be great enough to outweigh the chance of being destroyed.

Tradeoffs involving these factors may be critical in route selection. For example, Routes A and B (Figure A-5) differ in how they allocate risk between ingress and egress. Route A plays it safe on the ingress, detouring significantly to avoid Threat 6; but on egress it passes quite close to Threat 16. Route B takes a more direct path to the target than Route A, placing it in jeopardy from Threat 6, but leaving it with enough fuel on egress to avoid Threat 16. It might be that Route B is on the whole safer (i.e., has a lower total lethality); but Route A might be preferable, even so, because it affords a better chance at the target. According to this model, choice between Route A and Route B depends on how much chance of damaging the target is worth how much risk to own aircraft.

An aid which focuses exclusively on overall lethality (or survivability), ignores this tradeoff. On the other hand, there may be some reluctance,

both at the individual and organizational level, to make full use of a prescriptive model. For example, weighting of target versus aircraft values may seem inappropriate in a squadron route planning session. The solution adopted by PRP is to expose users to displays and messages in which the existence and import of such tradeoffs is made clear, when they occur, without the need for explicit quantitative modeling of all factors.

A variety of displays are available in ANALYZE to help users evaluate routes:

LETHALITY: This display shows how lethality increases as a function of the time spent on a route (Figure A-28). "Lethality" is the probability of having been killed on or before the indicated time; it is a cumulative measure which increases as time passes. (In terms of the decision tree model, the total lethality on a route is 1-p₁p₂.) The LETHALITY display, in essence, breaks total lethality down into components contributed by different segments of the route. The user can determine which parts of a route are most dangerous by observing where the slope of the curve rises most steeply. Independent consideration of ingress lethality and overall lethality is facilitated by an indicator ("T") which shows where on the curve the target is reached. In our example, Route A is more dangerous on egress than on ingress.

When two routes are selected, LETHALITY provides a comparison between them (Figure A-29). We can see that Route A is considerably safer than Route B on ingress, but is riskier than Route B on egress. The two routes are approximately equal in overall lethality, but Route A involves less chance of destruction before reaching the target.

EXPLAIN: These displays clarify the causes of lethality on each segment of a route (Figures A-30, A-31). Threats which contribute to a route's lethality are indicated at the appropriate point along the top of the lethality curve, associated with a symbol which indicates the threat's most

probable generic classification (i.e., surface-to-air missile, anti-air artillery, or radar). Messages placed below the curve indicate other sources of lethality. For example, Figure A-30 tells us that fuel problems increase lethality toward the end of Route A; Figure A-31 indicates that an overly predictable (i.e., straight line) segment of Route B is responsible for a portion of its lethality.

When two routes are selected, the EXPLAIN display operates somewhat differently. Instead of showing all the sources of lethality for both routes, it shows only those sources that are not shared by both (Figure A-32). In other words, it highlights the ways in which the two routes differ. In the example, Threats 2 and 6 affect Route B, but not Route A; conversely, Threats 4, 13, and 1 affect Route A, but not Route B. The two routes are also distinguished by the factors of fuel and predictability.

In addition, this EXPLAIN display carries a set of messages drawing attention to the tradeoffs that enter into the evaluation. These messages draw the user's attention explicitly to the difference between effectiveness against the target (where Route A is better) and overall lethality (where the two routes are equal).

A final message notifies users that Threat 6 ID plays a critical role in the choice between Route A and B, and that the evidence underlying the current assessment is somewhat incomplete. Thus, users are encouraged to examine that parameter more closely (using TOPIC and FOCUS), and to contribute any additional evidence they may have (using MODIFY).

The next three figures show what the user might find if he chooses to pursue the question of Threat 6 ID. Figure A-33 displays the locations of Routes A and B in relation to laydown contours for Threat 6; clearly, Route B is well inside the threat's capability area, while Route A is well outside. Recall, however, that these contours represent a probabilistic combination of the contours representing Threat 6's distinct ID

possibilities. If the user now selects LAYDOWN-ID, he may view the laydown contours separately for each possibility. Figures A-34 and A-35 show that Route B has trouble with Threat 6 only if the threat is identified as an SA-10; neither route is within the danger zone if Threat 6 is an SA-9. By selecting ID (Figure A-14), the user may see what the current assessment of ID probabilities is (about 45% chance of an SA-9, and 50% chance of an SA-10). If he checks SENSITIVITY (Figure A-22), he will learn that a very small change in ID probability might make Route B preferable. If we further suppose that the user has recent data (e.g., from a returning pilot) that points in the direction of an SA-9, he can use MODIFY to adjust the ID probabilities accordingly. He may then return to ANALYZE to observe the results of the change for his evaluation.

!SEGMENT and !THREAT are modules that facilitate an information search of this kind. They automatically take the user through a sequence of TOPIC/FOCUS and MODIFY displays, in their order of importance for the route evaluation problem at hand. Importance, in this case, means that: (a) route selection is sensitive to the value of a parameter; (b) current evidence is incomplete; and (c) users may have data not available to the computer. In our above example, the user would first be given the opportunity to adjust Threat 6 ID values; then he would be shown any other parameters with a potential impact on his choice between Routes A and B.

!THREAT and !SEGMENT differ in the way they organize the information to which a user is exposed. !THREAT would take the user directly to the threat parameter display to be modified (e.g., Figure A-14 for ID, with a MODIFY submenu). !SEGMENT, on the other hand, first displays the route segment whose design is affected by the threat; and only then produces the relevant threat parameter display. The user steps through these sequences by use of NEXT and PRIOR.

ANALYZE provides two other displays that clarify critical factors in route planning:

- FUEL: This display shows how remaining fuel is affected by time on a route (Figure A-36). Users can determine which portions of a route are expected to be most costly in terms of fuel, and how much fuel is expected to remain at each point. Route A, in this example, crosses the "critical" threshold (20% fuel remaining) in the last stages of the egress.
- ALERTS: This display provides an overview of critical points for decision and action along a route (Figure A-37). It has a dual use: it tells the route planner what the pilot on a given route would have to face; and it provides an advance set of alerts for the pilot himself regarding decision points and risky areas. In this example, the decision points shown for Route A include (a) the possibility of receiving information enroute pertaining to Threat 6 identity and (b) the amount of fuel remaining shortly after striking the target. (These are not the same branches as in the case where the user has explicitly decided to generate a contingency plan (Figure A-27). Here, even though the user has chosen Route A, the aid encourages flexibility in regard to potential choice points along the way.) The ALERTS display also shows areas along the route which are potentially critical, even though they are not explicit choice points: e.g., areas of terrain masking from a threat, and route portions where small deviations from the planned route may be especially costly.

ELIMINATE and SELECT are complementary methods by which the user can narrow down his working set of potential routes. They correspond to different heuristic choice strategies and ways of organizing information (viz., by attributes or by options, respectively). In ELIMINATE, the user can specify minimum and maximum allowable values on various evaluative attributes (e.g., total lethality, ingress lethality, final remaining fuel, length of a leg); all routes falling outside these cut-points will be dropped. An Advisory prompt might, however, warn users when a route is rejected which scores quite well on other attributes. SELECT, on the other hand, enables the user to designate the number of options n (from 1 to 4) that he wants to retain; it then performs an overall evaluation of each option, using the aid's internal prescriptive model, and retains only the best n choices.

5.0 CONCLUSIONS

5.1 Feasibility of Personalized Decision Aiding

The Personalized Route Planner embodies in concrete form many of the decision-aiding concepts and principles developed in Section 2.0. In so doing, it represents a partial, and preliminary, test of the feasibility of providing decision support which is at the same time personalized and prescriptively adequate.

Figure 6 summarizes some important features of PRP within the personalized decision-aiding framework. It shows which functions of PRP have potential use in each phase of decision making: i.e., situation assessment, option generation, and choice. And it analyzes each of the PRP functions into personalizing, channeling, and advisory prompting components at that phase.

What are the next steps? We conclude with brief comments on just two of many topics: the design of PRP inference mechanisms and the generalization of the personalized decision-aiding framework to other decision aids at RADC.

5.2 Prescriptive Mechanisms for PRP

According to the design laid out in Section 4.0, PRP will require three separable but interconnected prescriptive models:

- an inference model to relate uncertainty parameters at different levels (Figure 4);
- a knowledge-based expert system component for route generation;
 and
- a decision-analytic model for evaluation of generated routes (e.g., Figure 5).

- Shafer, G. A mathematical theory of evidence. Princeton, NJ: Princeton University Press, 1976.
- Shaklee, H., and Fischhoff, B. Strategies of information search in causal analysis. Memory and Cognition, 1982.
- Simon, H.A. Models of man, social and rational: Mathematical essays on rational and human behavior in a social setting. NY: Wiley, 1957.
- Slovic, P., and Fischhoff, B. On the psychology of experimental surprises. Journal of Experimental Psychology: Human Perception and Performance, 1977, 3, 544-551.
- Slovic, P., Fischhoff, B., and Lichtenstein, S. Accident probabilities and seat belt usage: A psychological perspective. Accident Analysis and Prevention, 1978, 10, 281-285.
- Streufert, S., and Streufert, S.C. Stress and the measurement of task performance. Decision making in complex tasks (Technical Report 3). Hershey, PA: The Milton S. Hershey Medical Center, Department of Behavioral Science, 1981.
- Svenson, O. Process descriptions of decision making. Organizational Behavior and Human Performance, 1979, 23, 86-112.
- Systems Control Technology, Inc. Route planning aid: Functional description (Draft Report). 31 October 1983.
- Tversky, A. Elimination by aspects: A theory of choice. *Psychological Review*, 1972, 79(4), 281-299.
- Tversky, A., and Kahneman, D. Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 1973, 4, 207-232.
- Tversky, A., and Kahneman, D. The framing of decisions and the psychology of choice. Science, 30 January 1981, 211, 453-458.
- Tversky, A., and Kahneman, D. Extensional vs. intuitive reasoning: The conjunction fallacy in probability judgment. Stanford, CA: Stanford University, Department of Psychology, June 1983.
- von Neumann, J., and Morgenstern, O. Theory of games and economic behavior (2nd Ed.). Princeton, NJ: Princeton University Press, 1947.
- Wason, P.C. On the failure to eliminate hypotheses in a conceptual task. Quarterly Journal of Experimental Psychology, 1960, 12, 129-140.
- Watson, S.R. Decision analysis as a replacement for cost-benefit analysis. European Journal of Operational Research, 1981, 7, 242-248.

- Miller, R.S. A framework for establishing the need for and value of C^2 programs. Monterey, CA: U.S. Naval Postgraduate School, June 1980.
- Nau, R. The duality of fuzzy probabilities (Draft). New Orleans, LA: Tulane University, A.B. Freeman School of Business, 1984.
- Neisser, U. Cognition and reality. San Francisco, CA: W.H. Freeman and Company, 1976.
- Newell, A. The heuristic of George Polya and its relation to artificial intelligence (Report No. CMU-CS-81-133). Pittsburgh, PA: Carnegie-Mellon University, Department of Computer Science, July 1981.
- Newell, A., and Simon, H.A. Human problem solving. Englewood Cliffs, NJ: Prentice-Hall, 1972.
- Payne, J.W. Information processing theory: Some concepts and methods applied to decision research. In T.S. Wallston (Ed.), Cognitive processes in choice and decision behavior. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc., 1981.
- Payne, J.W. Contingent decision behavior: A review and discussion of issues (82-1). Durham, NC: Duke University, Fugua School of Business, 1982.
- Pitz, G.F., Sachs, N.J., and Heerboth, J. Procedures for eliciting choices in the analysis of individual decisions. *Organizational Behavior and Human Performance*, 1980, 26(3), 396-408.
- Reiter, R. A logic for default reasoning. Artificial Intelligence, 1980, 13, 81-132.
- Riemenschneider, R.A., and Rockmore, A.J. Route planning aid: User's manual (PAR Report 83-106). New Hartford, NY: PAR Technology Corporation, 16 September 1983.
- Riemenschneider, R.A., Rockmore, A.J., and Wikman, T.A.A. Route planning aid: System specification (PAR Report 83-135). New Hartford, NY: PAR Technology Corporation, 31 October 1983.
- Savage, L.J. The foundations of statistics. NY: Wiley, 1954.
- Schum, D., DuCharme, W., and DePitts, K. Research on human multistage probabilistic inference processes. Organizational Behavior and Human Performance, 1973, 10, 318-348.
- Schum, D.A., and Martin, A.W. Assessing the probative value of evidence in various inference structures (Research Report 81-02). Houston, TX: Rice University, 1981.

- Kahneman, D., Slovic, P., and Tversky, A. (Eds.). Judgment under uncertainty: Heuristics and biases. Cambridge, England: Cambridge University Press, 1982.
- Kahneman, D., and Tversky, A. Subjective probability: A judgment of representativeness. Cognitive Psychology, 1972, 3, 430-454.
- Kahneman, D., and Tversky, A. Prospect theory: An analysis of decision under risk. *Econometrica*, March 1979, 47(2), 263-291.
- Kahneman, D., and Tversky, A. The psychology of preferences. Scientific American, January 1982, 246, 160-173.
- Koriat, A., Lichtenstein, S., and Fischhoff, B. Reasons for confidence. Journal of Experimental Psychology: Human Learning and Memory, 1980, 6, 107-118.
- Lamberti, D., and Wallace, W.A. Information portrayal and decision support for command and control: An experimental approach. Troy, NY: Rensselaer Polytechnic Institute, School of Management, September 1984.
- Larkin, J., McDermott, J., Simon, D.P., and Simon, H.A. Expert and novice performance in solving physics problems. *Science*, 1980, 208, 1335-1342.
- Leddo, J., Abelson, R.P., and Gross, P.H. Conjunctive explanations: When two reasons are better than one. *Journal of Personality and Social Psychology*, 1984, 47(5), 933-943.
- Lichtenstein, S., Fischhoff, B., and Phillips, L.D. Calibration of probabilities: The state of the art to 1980. In D. Kahneman, P. Slovic, and A. Tversky (Eds.), Judgment under uncertainty: Heuristics and biases. NY: Cambridge University Press, 1982.
- Lopes, L.L. Doing the impossible: A note on induction and the experience of randowness. Madison, WI: University of Wisconsin, Department of Psychology, undated.
- Lopes, L.L. Procedural debiasing (WHIPP #15). Madison, WI: University of Wisconsin, Department of Psychology, October 1982.
- McCarthy, J. Epistomological problems of artificial intelligence.

 Proceedings of the Fifth International Joint Conference on Artificial
 Intelligence. Cambridge, MA: Massachusetts Institute of Technology, 1977.
- McCarthy, J. Circumscription--A form of non-monotonic reasoning. Artificial Intelligence, 1980, 13(1,2), 27-39.
- McDermott, D., and Doyle, J. Non-monotonic Logic I. Artificial Intelligence, 1980, 13, 41-72.

- Fischhoff, B. Debiasing. In D. Kahneman, P. Slovic, and A. Tversky (Eds.), Judgment under uncertainty: Heuristics and biases, Cambridge, England: Cambridge University Press, 1982, 422-444.
- Fitts, P.M. (Ed.) Human engineering for an effective air-navigation and traffic-control systeml. Columbus, OH: Ohio State University Research Foundation, 1951.
- Foley, J.D., Wallace, V.L., and Chan, P. The human factors of graphic interaction: Tasks and techniques (Report GWU-IIST-81-3). Washington, DC: Institute for Information Science and Technology, January 1981.
- Gates, K.H., and Figgins, T.L. Decision aids for target aggregation: Task IV Subjective, empirical, and technical evaluation of selected decision aids (RADC-TR-84-108, Vol. IV of nineteen). PAR Technology Corporation, May 1984.
- Gentner, D., and Stevens, A.L. (Eds.). Mental models. Hillsdale, NJ: Lawrence Erlbaum Associates, 1983.
- Gardenfors, P., Hansson, B., and Sahlin; N. (Eds.), Evidentiary value: Philosophical, judicial, and psychological aspects of a theory. Lund, Sweden: C.W.K. Gleerups, 1983.
- Gardenfors, P., and Sahlin, N. Unreliable probabilities, risk taking and decision making. Synthese, 1982, 53.
- Hallden, S. Indiciemekanismer, Tidsskrift fur Rettsvitenskap, 1973, 86, 55-64.
- Hamm, R.M. Task conditions versus stable individual differences as determinants of experts' judgment policies (Report No. 249). Boulder, CO: Institute for Cognitive Science, Center for Research on Judgment and Policy, September 1983.
- Hogarth, R.M. Beyond discrete biases: Functional and dysfunctional aspects of judgmental heuristics. *Psychological Bulletin*, 1981, 90(2), 197-217.
- Huber, G.P. Cognitive style as a basis for MIS and DSS designs: Much ado about nothing? Madison, WI: University of Wisconsin, 1982.
- Johnson, E.J. Expertise and decision under uncertainty: Performance and process. Cambridge, MA: Massachusetts Institute of Technology, February 25, 1985.
- Kadane, J.B., and Lichtenstein, S. A subjectivist view of calibration (Report 82-6). Eugene, OR: Decision Research, 1982.

- (Technical Report 82-4). Falls Church, VA: Decision Science Consortium, Inc., July 1982.
- Cohen, M.S., and Brown, R.V. Decision support for attack submarine commanders (Technical Report 80-11). Falls Church, VA: Decision Science Consortium, Inc., October 1980.
- Cohen, M.S., Brown, R.V., Seaver, D.A., and Ulvila, J.W. Operability in attack submarine combat systems: An exploratory overview (Technical Report 82-2). Falls Church, VA: Decision Science Consortium, Inc., April 1982.
- Cohen, M.S., Schum, D.A., Freeling, A.N.S., and Chinnis, J.O., Jr. On the art and science of hedging: Alternative theories on the assessment, combination, and reporting of uncertainty in intelligence analysis (Technical Report 84-6). Falls Church, VA: Decision Science Consortium, Inc., August 1984.
- Cohen, M.S., Watson, S.R., and Barrett, E. Alternative theories of inference in expert systems for image analysis (Technical Report 85-1). Falls Church, VA: Decision Science Consortium, Inc., January 1985.
- Cohen, P.R. Heuristic reasoning about uncertainty: An artificial intelligence approach (Report No. STAN-CS-83-986). Stanford, CA: Stanford University, Department of Computer Science, August 1983.
- deFinetti, B. Foresight: Its logical laws, its subjective sources. English translation in H.E. Kybert, Jr., and H.E. Smokler (Eds.), Studies in subjective probability. NY: Wiley, 1964. (Original: 1937)
- Dawes, R.M. The mind, the model, and the task. In F. Restle et al. (Eds.), Cognitive theory. Hillsdale, NJ: Lawrence Erlbaum Assoc., 1975, 1, 119-130.
- Doyle, J. A truth maintenance system. Artificial Intelligence, 1979, 12(3), 231-272.
- Driver, M.J. and Mock, T.J. Human information processing, decision theory style, and accounting information systems. *Accounting Review*, 1976, 50, 490-508.
- Edwards, W. (Ed.). Revisions of opinions by men and man-machine systems. *IEEE Transactions on Human Factors in Electronics*, 1966, 7(1).
- Einhorn, H.J. Learning from experience and suboptimal rules in decision making. In T.S. Wallsten (Ed.), Cognitive processes in choice and decision behavior. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc., 1980.
- Einhorn, H.J., and Hogarth, R.M. Ambiguity and uncertainty in probabilistic inference. Chicago, IL: University of Chicago, Center for Decision Research, June 1984.

REFERENCES

- Adelman, L., and Crowley, J.O. Decision aids for target aggregation: Route planning aid (RPA) Test analysis (RADC-TR-84-108, Vol. XII (of nineteen). PAR Technology Corporation, May 1984.
- Beach, L.R., and Mitchell, T.R. A contingency model for the selection of decision strategies. Academy of Management Review, 1978, 3, 439-449.
- Beard, L.H. Planning a management information system: Some caveats and contemplations. Financial Executive, May 1977, 34-39.
- Brown, R.V., Kahr, A.S., and Peterson, C.R. Decision analysis for the manager. NY: Holt, Rinehart and Winston, 1974.
- Buchanan, B.G. Research on expert systems (Report No. HPP-8-1). Stanford, CA: Computer Science Department, School of Humanities and Sciences, February 1981.
- Buchanan, B.G., and Shortliffe, E.H. Rule-based expert systems. Addison-Wesley, 1984.
- Card, S.K., Moran, T.P., and Newell, A. The psychology of human-computer interaction. Hillsdale, NJ: Lawrence Erlbaum Associates, 1983.
- Chase, W. Elementary information processes. In W.K. Estes (Ed.), Handbook of learning and cognitive processes, Vol. 5. Hillsdale, NJ: Erlbaum, 1978.
- Chi, M., Feltovich, P., and Glaser, R. Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 1981, 5, 121-152.
- Chinnis, J.O., Jr., Cohen, M.S., and Bresnick, T.A. Human and computer task allocation in air-defense systems (Technical Report 84-2). Falls Church, VA: Decision Science Consortium, Inc., August 1984.
- Cohen, L.J. The probable and the provable. Oxford, England: Clarendon Press, 1977.
- Cohen, L.J. Can human irrationality be experimentally demonstrated? The Behavioral and Brain Sciences, 1981, 4(3), 317-330.
- Cohen, M.S. Decision support for attack submarine commanders: Attack planning and target range pooling (Technical Report 82-1). Falls Church, VA: Decision Science Consortium, Inc., April 1982.
- Cohen, M.S., Bromage, R.C., Chinnis, J.O., Jr., Payne, J.W., and Ulvila, J.W. A personalized and prescriptive attack planning decision aid

aids such as DART and RPA.

Interactive Methods. Some of the displays and interactive methods developed for the Personalized Route Planner may carry over almost directly to DART: in particular, the probability histogram used to represent ID uncertainty (Figure A-14), the LOCATION displays for OVERVIEW (Figure A-9), REGION, and THREAT (Figure A-10), and the MODIFY procedures for user adjustments of ID and LOCATION uncertainty (Figures A-20 and A-21).

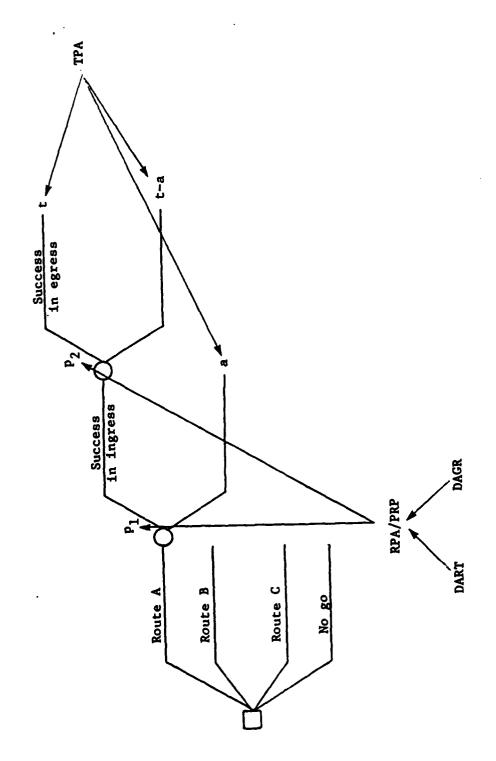
Some of the differences between DART and PRP, however, are also of interest. PRP is primarily concerned with choices among routes; threat classification is only one among many relevant inputs. Thus, PRP leaves inferences about threat ID primarily up to the computer, occasionally prompting the user when a user contribution might be of value (see Section 2.6.2 above). DART, on the other hand, focuses primarily on the inference task. Hence, a change in the balance of initiative may be appropriate, with an expectation of greater user involvement. At the same time, Advisory prompts might be designed to counterbalance this involvement, warning users, for example, when uncertainty appears to be suppressed (e.g., in conjunctive hypotheses or long chains of reasoning) or when disconfirming evidence has been ignored (Section 2.6.1.1).

A variety of other PRP displays may be of considerable use to higher-level mission planners in the TACC. TOPIC/FOCUS displays provide a good overview of a mission area and the implications of a particular route. LETHALITY and EXPLAIN displays might help planners balance benefits and costs in selecting targets and in allocating resources to missions.

Unfortunately, as they now stand, the Decision Aids for Target Aggregation (DATA) are neither consistent with one another nor entirely adequate in their own right. Potential improvements include both the underlying uncertainty models and the interactive methods by which results are communicated to users.

Uncertainty Models. Measures of uncertainty in DART are acknowledged to be highly ad hoc, and do not lend themselves readily to any clear interpretation. The problem, however, will not be solved simply by designing clearer displays, or by finding translations of DART's uncertainty measures in terms of natural expressions like "very certain". Rather, some of the basic methods for manipulating uncertainty in DART are simply not justifiable; hence, the outputs may be meaningless, regardless of how they are displayed or translated. For example, when a hypothesis regarding the ID of a threat is added to DART's database, competing hypotheses are dropped, even if the difference in support between the favored hypotheses and the other possibilities is very small. (This is an example of a cognitive bias, discussed in Section 2.5), in which one behaves "as if" an uncertain conclusion were known for sure.) Further, new evidence which is incompatible with an established classification is not utilized by DART. But such new evidence (if taken seriously) would rule out the currently favored ID hypothesis and support a hypothesis that was previously dropped. Thus, there is no assurance that uncertainty measures in DART will produce a sensible rank ordering of possible classifications.

A variety of inference mechanisms might be applied to the DART problem; several are reviewed by Cohen, Watson, and Barrett (1985): e.g., Bayesian probabilities or Shaferian belief functions. The Bayesian system is well understood and has a highly plausible, if not compelling, normative justification. The Shaferian system seems to be better at capturing intuitions about confirmation and about incompleteness of evidence (Section 2.6.4 above). On the whole, we suspect that a Bayesian representation may provide a sound, consistent framework for reasoning about uncertainty in



Possible sources of inputs for the route planning decision tree. Figure 7:

An expert system route generator might utilize a mix of methods. One approach is to start with the "minimum lethality" route (as presently computed) and allow heuristic rules to modify it to accommodate omitted factors. Another method is to start with a "minimum fuel" route (essentially a set of straight lines linking user-provided constraints), and modify it to take account of specific threats. (The latter resembles the "object oriented" approach currently being explored at RADC.) Applying multiple approaches is one way to arrive at qualitatively different route options for evaluation.

Both the inference model and the decision tree model could be developed by straightforward applications of probability theory and decision analysis. The route generator will itself make use of outputs from each of these two sources.

5.3 Generalization to other RADC Aids

Ideas developed in the design of PRP could be extended, in various ways, to other RADC aids: e.g., DART, DAGR, TPA, and others. The interrelations among the functions of these aids (and the possible overlap of users) argues the need for a consistent framework of some sort.

In principle, the route planning process can, and should, make use of outputs from these aids. Figure 7 shows a decision tree for route selection, in which probabilities for outcomes (p_1,p_2) and values for those outcomes (t,a) are derived from other aids. Thus, probability of survival on a route is derived from evidence regarding the classification of threat C^2 installations (DART) and regarding the enemy order of battle (DAGR). The relative value of target and own forces may be based on evaluative criteria included in TPA. It seems clear that a consistent prescriptive framework for handling uncertainty and balancing competing objectives would be of value for this group of aids.

The substitution of an expert system route generator for the present optimization algorithm is central to the recommendations underlying the PRP design. Many of these features could not be implemented with the present algorithm; for example:

- the incorporation of fuzzy and partial user constraints;
- the generation of an entire route (versus ingress and egress separately);
- the generation of multiple candidate routes which are qualitatively different;
- the incorporation of branches or contingency plans;
- the ability to segment routes into portions whose design was influenced by particular threats; and
- the influence of global attributes, like fuel flow and avoiding legs that are too long or too short.

Our suggestion is that these features, which define the user's interaction with the aid, should serve as constraints in the design of an expert system model, rather than the other way around. It can be a serious error to let the capabilities of a particular model determine the way users interact with the aid. But to a large degree, this is what has happened with RPA. For example, the present algorithm can only generate routes connecting two points. As a result, users must commit themselves, however artificially, to exact ingress and egress points, and cannot utilize other constraints they might find desirable or useful; for the same reason, routes must be generated piecemeal rather than whole. Further, the essentially local operations performed by that algorithm are very different from a human's natural way of conceptualizing routes: as a result, it is meaningless to ask about alternative routes, contingencies, or the impact of a particular threat. Finally, users are left to supply judgments on a variety of topics (fuel, predictability, etc.) not because humans are particularly suited for this task, but because the aid cannot deal with it at all.

PROMPTING		Prompt when incompleteness of evidence regarding a factor affects choice			Guide user through uncertainty factors in order of incomplete- ness of evidence and impact on choice	Prompt if strategy misses potentially desirable options
CHANNELING	facilitates balancing tradeoffs (fuel, danger; effectiveness, lethality)		Highlight segment and threat as decision-related objects	Highlight important sources of uncertainty and inferential relationships Exploit user knowledge where it matters for choice	Focus attention on decision-related objects (segment, threat)	
PERSONAL 121NG	Show lethality for total route and by segment Prescriptive model not made explicit	Show causal factors affecting route lethality Explain differences between routes by "advantages" and "disadvantages"	View route information according to meaningful "objects" (segment, threat, leg, region)	View impact on route of uncertainty information at any level View impact of user values (at any level) on lethality	Organize information search by different objects	facilitate alternative choice strategies based on differ- ent organization of knowledge
AID FEATURE	ANALYZE (LETHALITY)	ANALYZE (EXPLAIN)	Focus	TOPIC MODIFY (SENSITIVITY)	ANALYZE (!THREAT, !SEGMENT)	ANALYZE (ELIMINATE, SELECT)

OPTION GENERATION

PROMFTING	Prompt when incompleterness of evidence makes a contingency plan appropriate	Prompt if piecewise generation of route is suboptimal	Prompt when incompleteness of evidence about specific factors affects viable option set	Prompt when a contingency plan might be appropriate	70	ě
CHANNELING	Facilitate consideration of multiple options	Facilitate appropriate time horizon (total route)	facilitate consideration of multiple factors, balance of ingress and egress	Facilitate view of option as a contingency plan (vs. fixed procedure)	Give pilot information required for decisions	Facilitates view of route as one part of larger option package
PERSONAL 121NG	Permit fuzzy inputs and partial constraints, blend of automatic and manual	Any portion of route can be generated in any order	Generate routes in relation to any meaningful object (threat, region) at any level in uncertainty pyramid	Meaningful partition of route into decision-based segments	Show fuzziness of leeway along route corridor	Ability to modify target, tot, a/c, etc. at any time
AID FEATURE	GENERATE	GENERATE	FOCUS, TOPIC	ANALYZE (ALERTS)		CONSTRAINTS

PROMPTING

CHANNELING	
PERSONAL 12 1NG	Organize information by different "objects" (threat, region)
AID FEATURE	Focus

inferen-	user con- itedge base; /alues as
Nighlight important sources of uncertainty and inferential relationships	Exploit potential user contributions to knowledge base; provide original values as reference points
View information at any level in inferential "pyramid"	User can insert own values at any level, permanently or for "what-if" analysis
TOPIC	MODIFY

Figure 6: PRP functions as seen in the personalized decision aiding framework.

Watson, S.R., Weiss, J.J., and Donnell, M.L. Fuzzy decision analysis. IEEE Transactions of Systems, Man and Cybernetics, 1979, SMC-9, 1-9.

Zadeh, L.A. Fuzzy sets. Inf. and Contr., 1965, 8, 338-353.

Zadeh, L.A. The concept of a linguistic variable and its application to approximate reasoning. *Information Science*, 1975, 8, 199-249; 301-357; 9, 43-80.

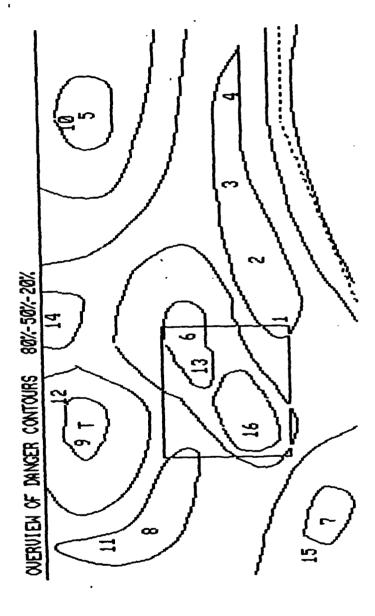
Zadeh, L.A. The role of fuzzy logic in the management of uncertainty in expert systems. Fuzzy Sets and Systems, 1983, 11, 199-227.

APPENDIX: DISPLAY SCREENS

ATO Rec'd	:	Control pt's 1	2	····	4	Aim pt	Ordnance =	T.0.T. =		
Display Lat to coord's Lon to	ŀ		7.	3	3	4/C = F111 F14 F7	Jamming (F1114 ECM) = Yes No	Visibility = Poor((1 nm) Good()1 nm)	Climate: Summer Winter Weight: hi med to	•

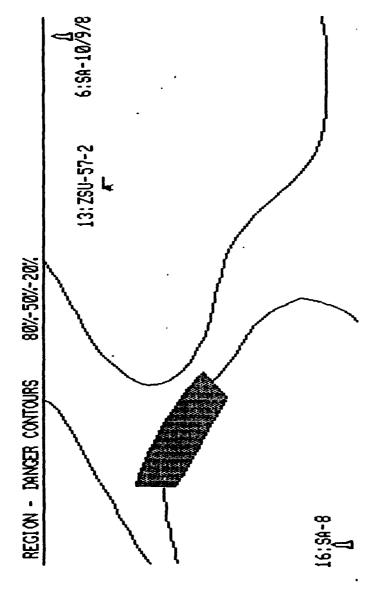
CONSTITUTES FOCUS TOPIC MODIFY GENERATE ANALYZE REPORTS QUIT Specify target, aircraft, weather, etc.

Figure A-1



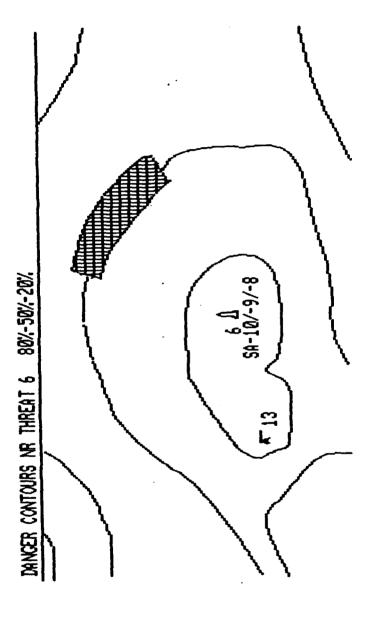
OVERVIEW REGION THREAT ROUTE SEGMENT LEG GRID Select a region with the mouse.

Figure A-2



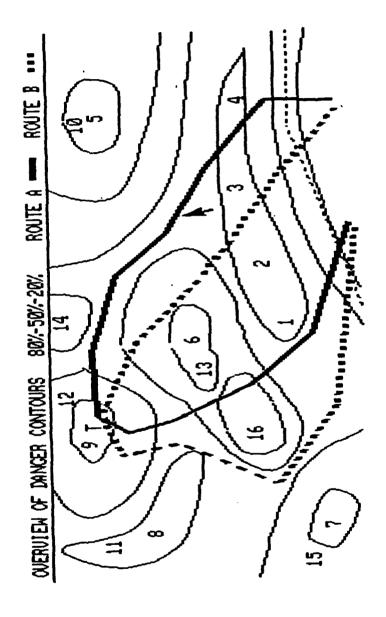
OUERVIEW WANTON THREAT ROUTE SECMENT LEG GRID Select a region with the mouse.

Figure A-3



OVERVIEW REGION WHATH ROUTE SEGMENT LEG CRID Select a threat with the mouse.

Figure A-4



OUERUIEM REGION THREAT ROUTE SEGMENT LEG GRID Select one or more routes with the mouse.

Figure A-5

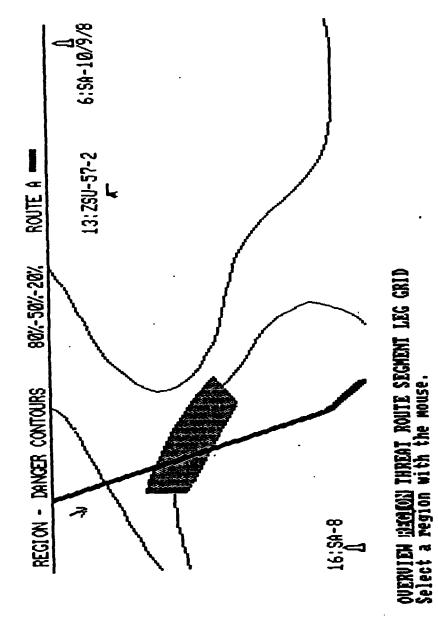
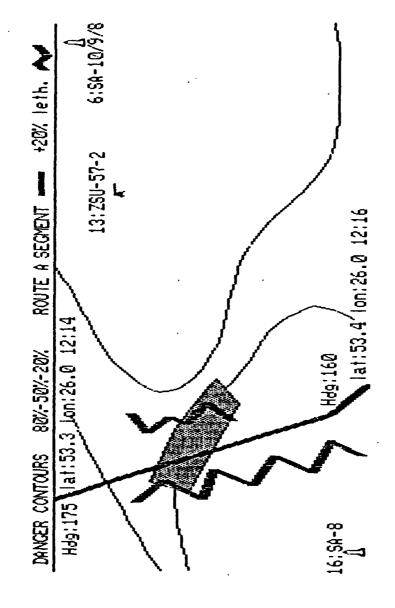
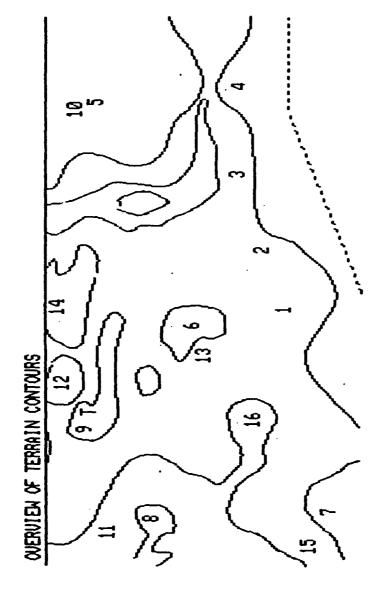


Figure A-6



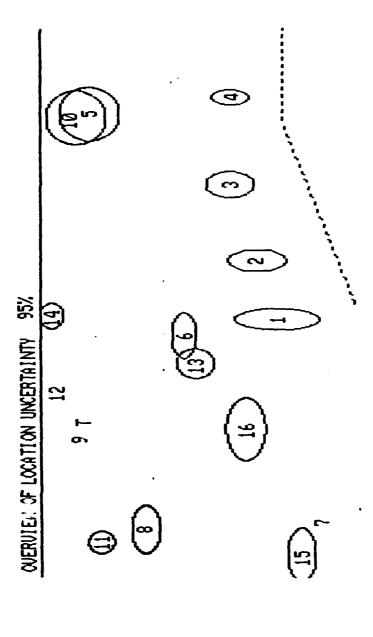
OUERUIEM REGION THREAT ROUTE SAGINING LEG GRID Select a segment of a route with the mouse.

Figure A-7



MUNICIPAL REGION THREAT ROUTE SECNENT LEG GRID DISPLAY the entire area.

Figure A-8



DANGER LAYDOWN LAYDOWN-ID ID MOSTREON TERRAIN CAPABILITY THREAT-TYPE Show probable areas of threat location.

Figure A-9

LOCATION UNCERT, FOR THREAT 6	UNCERT.	æ	TREAT		95%			
							•	•
•	•			•	•	•		
***************************************		:						
	• •					•		
•	•••••••••••••••••••••••••••••••••••••••	:	:	: :				•
•	•			•	•		• .	•
	•	:	::			:.		
	•						•	
						,		,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,
•	•		•			•	•	•
	-		•	.]	i	-	-	
	•		`		•	-	-	•
	• •		<u> </u>	•			•	-
		:::::::	<i>)</i>		· · · · · · · · · · · · · · · · · · ·			[
•	•		•	k			•	
	-	-	•	٠,	0_//_	-		
		<u>-</u>						
	•	•	•	•	•	•	•	
								•
•	•			•	•	•	-	•
	•		•	- 1	•	•	• .	
	.,		•••••••••••••••••••••••••••••••••••••••	•••••••••••••••••••••••••••••••••••••••			••••••••	
-	•				•			
	•		•	•				

DANGER LAYDOWN LAYDOWN-ID ID MOGNATOR TERRAIN CAPABILITY THREAT-TYPE Show probable areas of threat location.

Figure A-10

50% DANGER CONTOURS

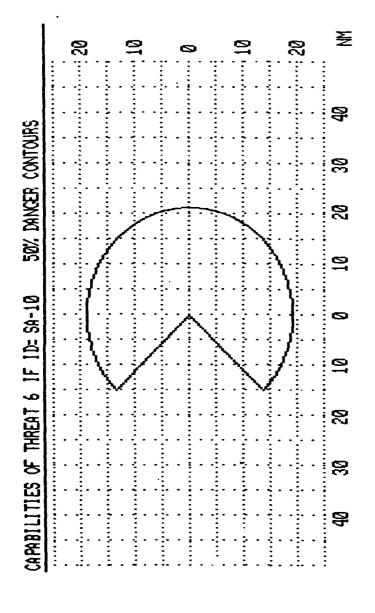
CAPABILITIES OF THREAT 6 IF ID: SA-9

₹

8

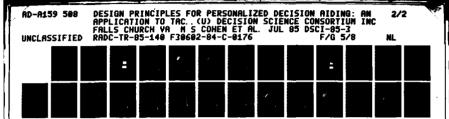
SA-10 STES SA-8 COMBINED Assume the threat is a SA-9.

Figure A-11

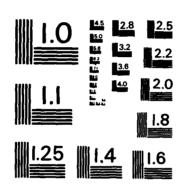


SECTION SA-9 SA-8 COMBINED ASSUME the threat is a SA-10.

Figure A-12

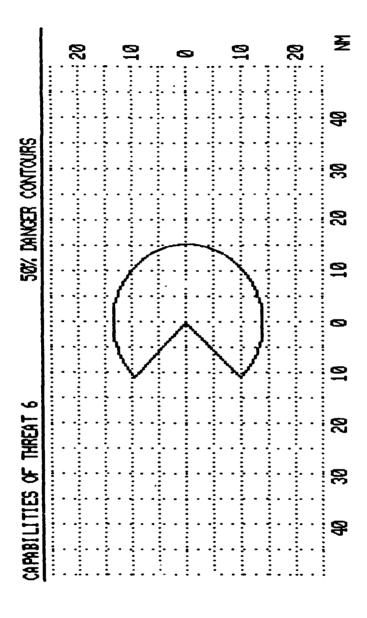


END FILMED OTIC



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

The second of the second secon



SA-18 SA-9 SA-8 303341343 Show threat capability combining all potential IDs.

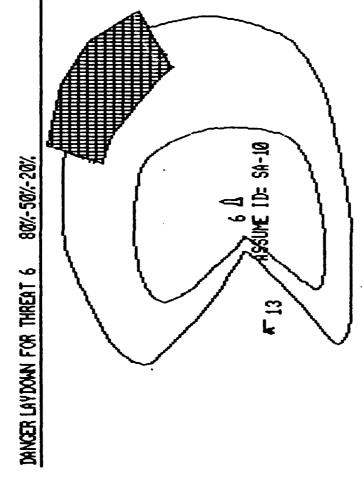
Figure A-13

Figure A-14

errord consisted rezerved ferencest research forested between between

DANCER LAYDOWN FOR THREAT 6 80%-50%-20%

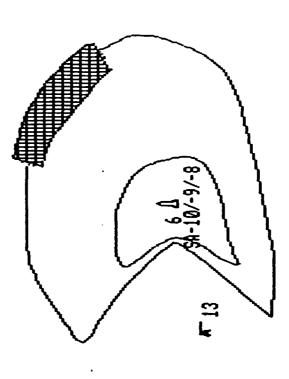
SA-18 STEP SA-8 Assume the threat is a SA-9. Figure A-15



SINIC SA-9 SA-8 Assume the threat is a SA-18.

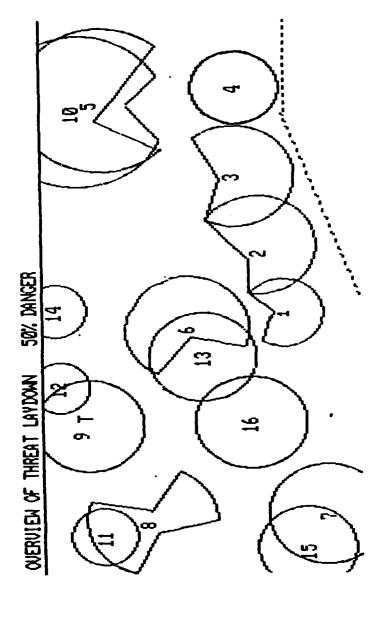
Figure A-16

del sesserel sessessi frances



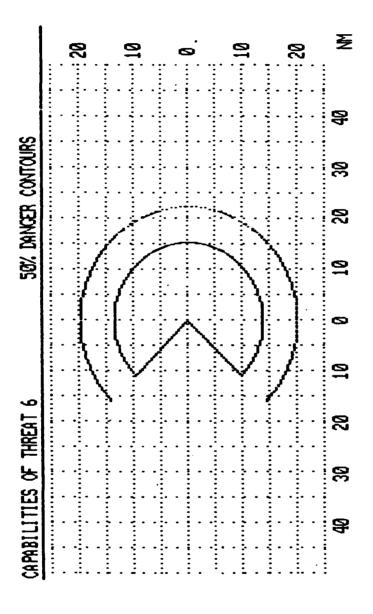
DANCER MINION LAYDOWN-ID ID LOCATION TERRAIN CAPABILITY THREAT-TYPE Show contours reflecting probability of kill from each threat.

Figure A-17



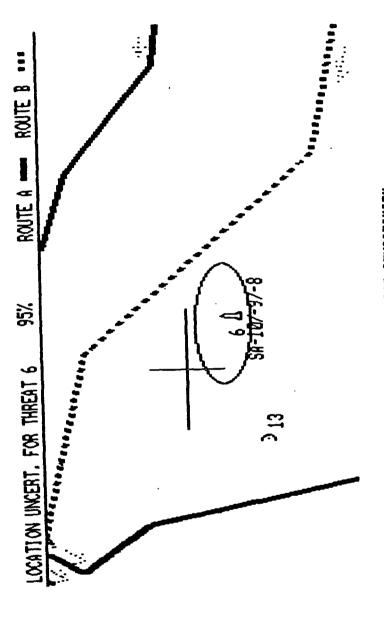
DANGER INVIOUS LAYDOWN-ID ID LOCATION TERRAIN CAPABILITY THREAT-TYPE Show contours reflecting probability of kill from each threat.

Figure A-18



RANCE SPREAD ORIENTATION UNDO DEFAULT EXPLAIN NEXT PRIOR SENSITIUITY Adjust the range of the threat.

Figure A-19



POSITION AXES UNDO DEFAULT EXPLAIN NEXT PRIOR SENSITIUITY Adjust the x-y position of the location ellipse.

Figure A-20

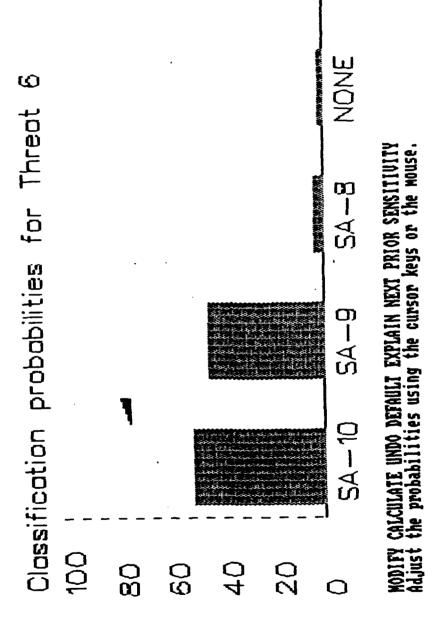
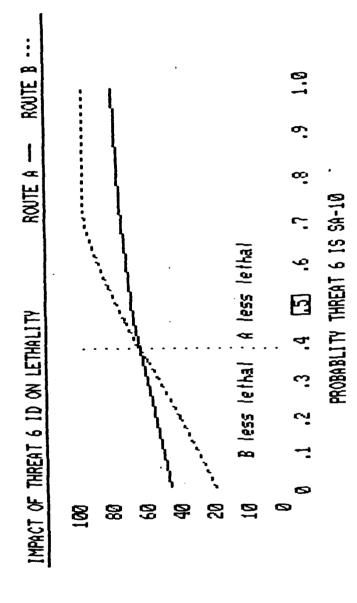
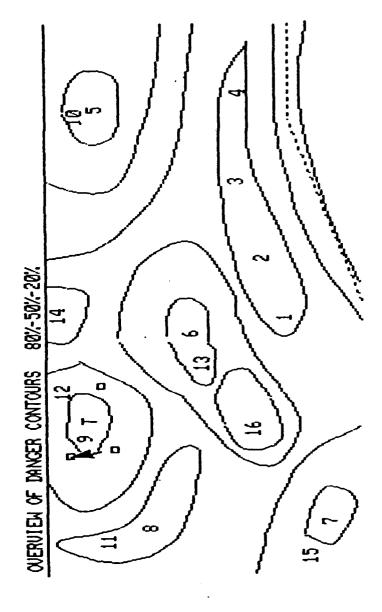


Figure A-21



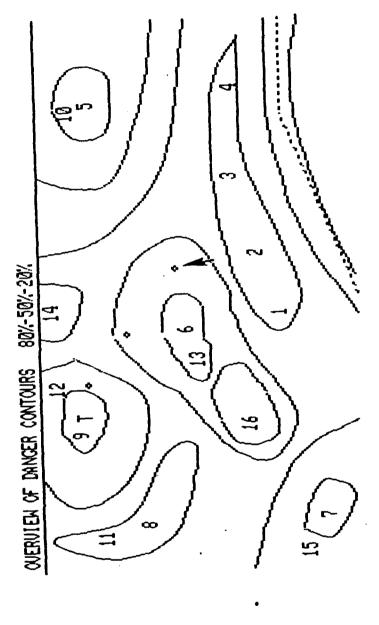
MODIFY CALCULATE UNDO DEFAULT EXPLAIN NEXT PRIOR SANSIBLUSES. Show impact on route lethality of changes in probability.

Figure A-22



ATTACK INGRESS EGRESS POINTS DRAW ALT/SPEED AUTO CONTINGENCY Draw or erase possible IPs with the mouse.

Figure A-23



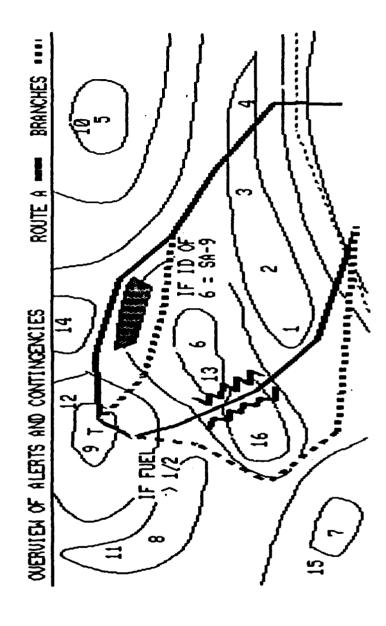
ATTACK INGRESS EGRESS POINTS DRAW ALT/SPEED AUTO CONTINGENCY Draw or erase route through-points with the mouse.

Figure A-24

THE CONTROL OF THE CO

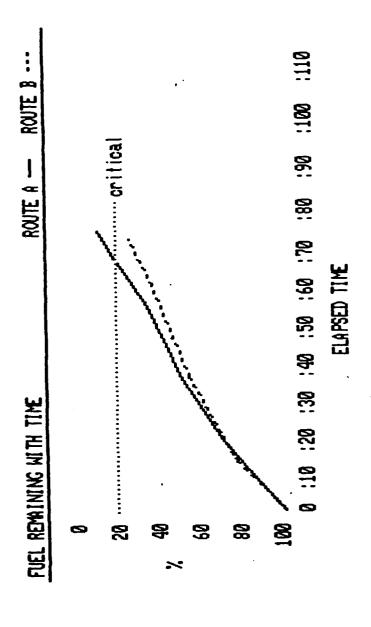
MISSION of Rome Air Development Center

RADC plans and executes research, development, test and selected acquisition programs in support of Command, Control Communications and Intelligence (C³I) activities. Technical and engineering support within areas of technical competence is provided to ESD Program Offices (POs) and other ESD elements. The principal technical mission areas are communications, electromagnetic guidance and control, surveillance of ground and aerospace objects, intelligence data collection and handling, information system technology, ionospheric propagation, solid state sciences, microwave physics and electronic reliability, maintainability and compatibility.



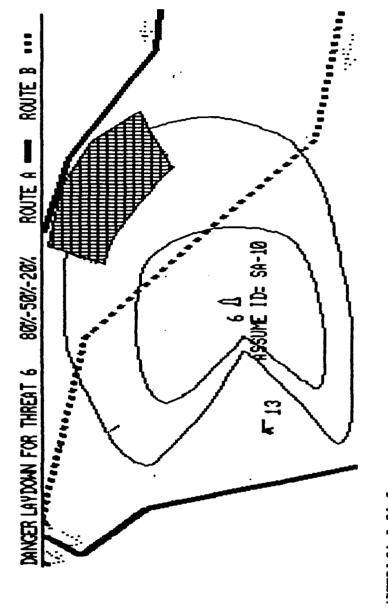
LETHALITY EXPLAIN FUEL MANAS ! SECHENT ! THREAT ELIMINATE SELECT Show decision-making points along route.

Figure A-37



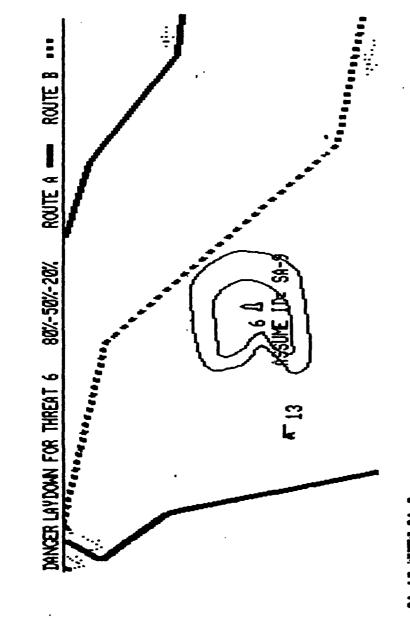
LETHALITY EXPLAIN <u>WHAN</u> ALERTS !SECHENT !THREAT ELIMINATE SELECT Show route fuel flow.

Figure A-36



SINIC SA-9 SA-8 Assume the threat is a SA-10.

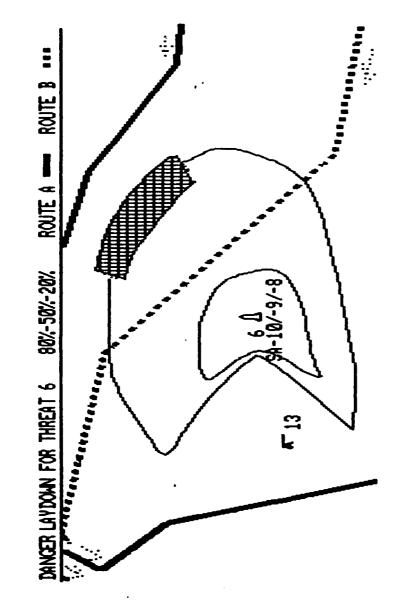
Figure A-35



SA-18 3752 SA-8 Assume the threat is a SA-9.

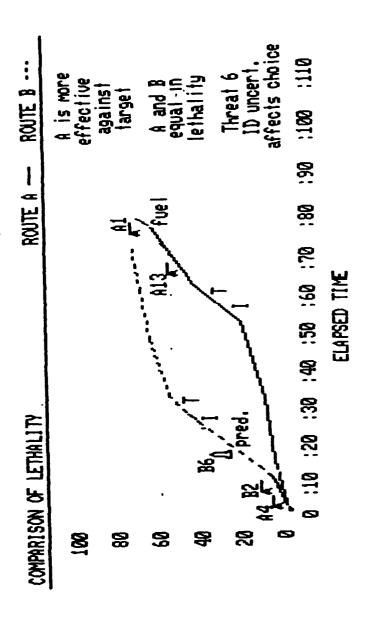
Figure A-34

Makerera desistan bandeler bandeler benerata desista



DANGER ICHICALLY LAYDOWN-ID ID LOCATION TERRAIN CAPABILITY THREAT-TYPE Show contours reflecting probability of kill from each threat.

Figure A-33



LETHALITY MANAGORI FUEL ALERTS ISECNENT ITHREAT ELIMINATE SELECT Indicate major contributors to route lethality.

Figure A-32

AND KREEKS BARROW BARR

CAUSES OF LETHALITY 9 鲁

LETHALITY MANAGIN FUEL ALERTS SECHENT STHREAT ELIMINATE SELECT Indicate major contributors to route lethality.

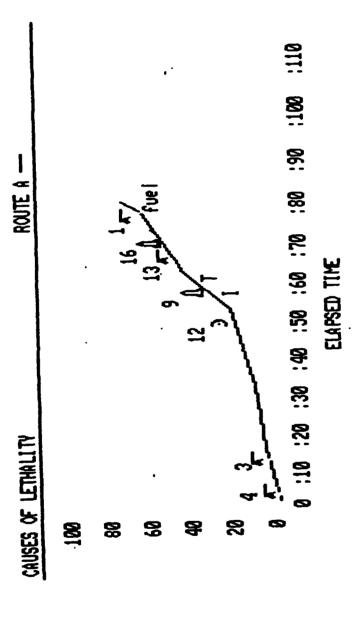
:110

:180

06 : 10 : 20 : 30 : 40 : 50 : 60 : 70 : 80 : 90

ELAPSED TIME

Figure A-31



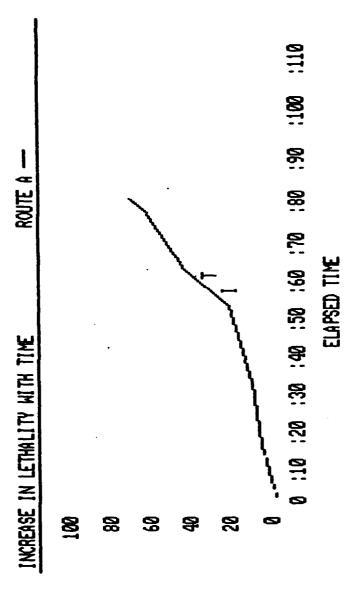
LETHALITY MANAGEM FUEL ALERTS ISECHENT ITHREAT ELIMINATE SELECT Indicate major contributors to route lethality.

Figure A-30



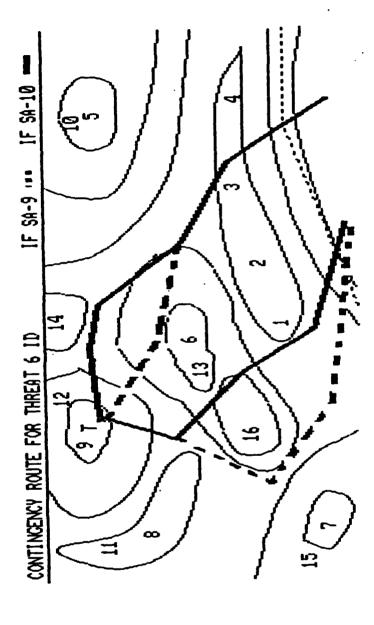
patishees explain fuel alerts iscenent ithreat eliminate select show how route lethality increases with time.

Figure A-29



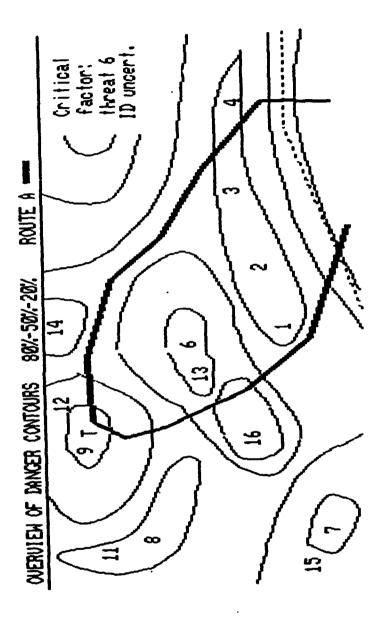
INVIENDED EXPLAIN FUEL ALERTS SEGMENT STHREAT ELIMINATE SELECT Show how route lethality increases with time.

Figure A-28



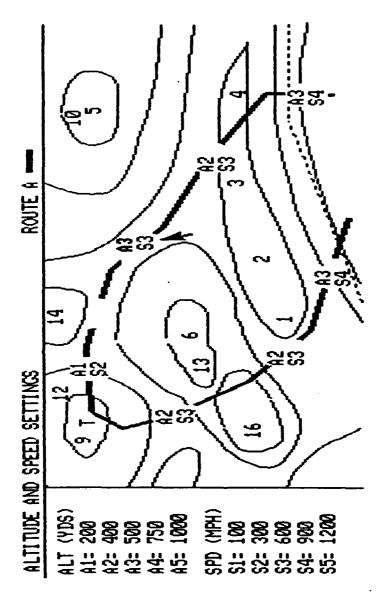
ATTACK INGRESS EGRESS POINTS DRAW ALT/SPEED AUTO WANDINGAINS! Generate branches where enroute data would affect choice.

Figure A-27



ig 2 3 4 Display only the best route.

Figure A-26



ATTACK INCRESS EGRESS POINTS DRAW <u>INTRIBUTATION</u> AUTO CONTINCENCY Show or specify altitude and speed constraints.

Figure A-25

END

FILMED

11-85

DTIC